



Firewall

FortiOS™ Handbook v3
for FortiOS 4.0 MR3



FortiOS™ Handbook Firewall

v3

24 January 2012

01-432-148222-20120124

© Copyright 2012 Fortinet, Inc. All rights reserved. Contents and terms are subject to change by Fortinet without prior notice. Reproduction or transmission of this publication is encouraged.

Trademarks

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Visit these links for more information and documentation for your Fortinet products:


Fortinet Knowledge Base - <http://kb.fortinet.com>

Technical Documentation - <http://docs.fortinet.com>

Training Services - <http://campus.training.fortinet.com>

Technical Support - <http://support.fortinet.com>

You can report errors or omissions in this or any Fortinet technical document to techdoc@fortinet.com.



Contents

Introduction	9
Before you begin	9
How this guide is organized	9
Understanding the FortiGate firewall	11
What is the FortiGate firewall?	11
FortiGate firewall components	11
How the firewall components create a FortiGate firewall and help in protecting your network.	12
Understanding how a packet travels through the FortiGate unit.	13
How packets flow in and out of the FortiGate unit.	14
Working with NAT in FortiOS	17
NAT in FortiOS	17
NAT/Route mode.	17
Route mode	18
Transparent mode	19
Types of NAT in FortiOS	19
Static NAT (SNAT)	20
Static Destination NAT (SDNAT).	20
Static NAT port forwarding	20
Dynamic NAT (DNAT)	20
Dynamic source address translation.	21
Dynamic destination address	21
Dynamic port forwarding	21
Combining types of NAT	21
Firewall components	23
Using Interfaces and zones in the FortiGate firewall	23
How to apply VLANs and zones and to a security policy	23
Understanding the firewall address component	24
IP addresses for self-originated traffic	25
IP pools.	26
IP Pools for security policies that use fixed ports	27
Source IP address and IP pool address matching.	27
Geography-based addressing	28
Wildcard addresses	29

Using wildcard addresses in the firewall configuration	31
Fully Qualified Domain Name addresses	31
Address groups	32
Virtual IP addresses	32
Grouping virtual IPs	32
Match-vip	32
Services	33
Predefined service list	33
Service groups	40
Firewall schedules	41
Schedule groups	41
Schedule expiry	41
UTM profiles	42
How to use UTM profiles to monitor and protect your network	42
Security policies	45
Security policy overview	45
Security policy list details	46
Viewing security policies	47
Policy order	47
How to arrange policies	49
Security policies	49
Identity-based policies	50
Identity-based policy example	51
SSL VPN policies	51
IPsec policies	52
Accept policies	52
Deny policies	52
How to allow DNS queries to only one DNS server	52
IPv6 policies	53
Security policy 0	53
Local-in policies	54
Creating a basic security policy	54
How to create a basic security policy for Internet access	55
How to test the basic security policy	55
How to verify if traffic is hitting the basic security policy	56
Monitoring firewall traffic	57
Session tables	57
Viewing session tables in the web-based manager	57
Sessions Monitor	57
Viewing session tables in the CLI	58
Proto_state fields: TCP	59
Proto_state fields: SCTP	60

Proto_state fields: UDP	60
Proto_state field for ICMP	60
Monitoring security policy traffic activity	60
Internet Protocol version 6 (IPv6)	63
What is IPv6?	63
IPv6 in FortiOS	64
Dual stack routing configuration	64
IPv4 tunneling configuration	65
Remotely connecting to an IPv6 network over the Internet	65
IPv6 overview.	65
Differences between IPv4 and IPv6	65
IPv6 MTU.	66
IPv6 address format	66
IP address notation	67
Netmasks.	68
Address scopes	68
Address types	68
Unicast.	68
Multicast	68
Anycast	69
Special addresses	69
Header Extension	70
IPv6 neighbor discovery	72
Transition from IPv4 to IPv6	73
Configuring FortiOS to connect to an IPv6 tunnel provider	74
Create a SIT-tunnel interface.	75
Create a static IPv6 route into the tunnel-Interface	75
Assign your IPv6 network to your FortiGate	75
Create a security policy to allow traffic from port1 to the tunnel interface	76
Test the connection	76
FortiGate IPv6 configuration	76
Displaying IPv6 options on the web-based manager	77
UTM protection for IPv6 networks	77
Configuring IPv6 interfaces	77
IPv6 interfaces - web-based manager.	77
IPv6 interfaces - CLI	77
Configuring IPv6 routing	78
Static routing.	78
Dynamic routing	79
Configuring IPv6 security policies	79


IPv6 Policy configuration settings	79
Configuring IPv6 DNS	83
Configuring IPv6 DHCP	83
Configuring IPv6 over IPv4 tunneling	83
Configuring IPv6 IPsec VPNs	84
Certificates	85
Configuring IPv6 IPsec VPNs	85
Security policies	86
Routing	86
IPv6 troubleshooting	86
ping6	87
IPv6 ping description	88
IPv6 ping options	88
Examples	90
diagnose sniffer packet	90
diagnose debug flow	91
IPv6 specific diag commands	91
Additional IPv6 resources	91
Advanced FortiGate firewall concepts	93
Central NAT table	93
Central NAT Table configuration settings	94
Stateful inspection of SCTP traffic	94
Configuring FortiGate SCTP filtering	95
Adding an SCTP custom service	96
Adding an SCTP policy route	96
Changing the session time to live for SCTP traffic	97
Port pairing	97
Blocking port 25 to email server traffic	98
Dedicated traffic	99
Restricting traffic on port 25	100
Blocking HTTP access by IP	101
ICMP packet processing	102
Adding NAT security policies in Transparent mode	102
Adding a static NAT virtual IP for a single IP address and port	105
Double NAT: combining IP pool with virtual IP	107
Using VIP range for Source NAT (SNAT) and static 1-to-1 mapping	109
Traffic shaping and per-IP traffic shaping	111
Endpoint Security	112
Logging traffic	112
Quality of Service (QoS)	113

Identity-based security policies 113

 Identity-based policy positioning 113

 Identity-based sub-policies 114

Index 115



Introduction

Welcome and thank you for selecting Fortinet products for your network protection. This document describes how to configure the FortiGate firewall on your FortiGate unit. This document also provides advanced firewall concepts.

This chapter contains the following topics:

- [Before you begin](#)
- [How this guide is organized](#)

Before you begin

Before you begin using this guide, please ensure that:

- You have administrative access to the web-based manager and/or CLI.
- The FortiGate unit is integrated into your network.
- The operation mode has been configured.
- The system time, DNS settings, administrator password, and network interfaces have been configured.
- Firmware, FortiGuard Antivirus and FortiGuard Antispam updates are completed.
- FortiGuard Analysis & Management Service is properly configured.

While using the instructions in this guide, note that administrators are assumed to be super_admin administrators unless otherwise specified. Some restrictions will apply to other administrators.

How this guide is organized

This FortiOS Handbook chapter contains the following sections:

[Understanding the FortiGate firewall](#) provides general information about what the FortiGate firewall does, what it is comprised of, and explains how a packet travels through the FortiGate unit.

[Working with NAT in FortiOS](#) provides information about how NAT works in FortiOS and the combinations of NAT that you can use in your configuration. This section explains how the different modes, such as Transparent mode, work and how the FortiGate unit behaves when in each of these modes.


[Firewall components](#) provides in-depth information about the firewall components that help in creating a FortiGate firewall configuration.

[Security policies](#) explains what security policies are, as well as how these rules work to help protect your network. This section also explains the importance of how security policies are ordered within the security policy list, and describes the different policies that can be created for different firewall configurations.

[Monitoring firewall traffic](#) explains how you can monitor traffic within the web-based manager using the Session and Policy Monitoring pages.

[Internet Protocol version 6 \(IPv6\)](#) explains how IPv6 can be implemented in FortiOS, as well as what features support IPv6, such as IPsec VPN and dynamic routing. This section also explains a high-level summary of IPv6.

[Advanced FortiGate firewall concepts](#) explains the advanced firewall features that you may want to configure for your network, as it expands. This section explains advanced firewall features that include stateful inspection of SCTP traffic, port pairing (Transparent mode only), and adding NAT security policies in Transparent mode.



Understanding the FortiGate firewall

The FortiGate firewall is one of the most important features on the FortiGate unit, allowing not only traffic to flow through, but also, with the help of security policies, scan the traffic for vulnerabilities and misuse and abuse. This type of firewall provides flexibility for expansion in a growing network environment.

This section helps to explain the FortiGate firewall and its role in protecting your network. This section also explains the life of a packet, which helps you to understand how the traffic flows through the FortiGate unit and the role the FortiGate firewall plays in the life of a packet.

The following topics are included in this section:

- [What is the FortiGate firewall?](#)
- [FortiGate firewall components](#)
- [Understanding how a packet travels through the FortiGate unit](#)

What is the FortiGate firewall?

A firewall is, in the simplest of terms, a device that permits or denies network traffic based on a set of rules. For the FortiGate firewall, it can do this and much more. The FortiGate firewall scans the network traffic, and based on the set of rules (in Fortinet, however, these rules are called security policies), determines what action needs to be taken. The action may be to quarantine a virus that the FortiGate unit finds, or to record the activity, or both. These security policies provide the information the FortiGate unit needs to determine what to do with the incoming and outgoing traffic.

At the heart of these networking security functions, is the security policies. Security policies control all traffic attempting to pass through the FortiGate unit, and between FortiGate interfaces, zones, and VLAN subinterfaces. They are instructions the FortiGate unit uses to decide connection acceptance and packet processing for traffic attempting to pass through. When the firewall receives a connection packet, it analyzes the packet's source address, destination address, and service (by port number), and attempts to locate a security policy matching the packet.

Security policies can contain many instructions for the FortiGate unit to follow when it receives matching packets. Some instructions are required, such as whether to drop or accept and process the packets, while other instructions, such as logging and authentication, are optional. It is through these policies that the FortiGate unit permits or denies the packets to pass through to the network, who gets priority (bandwidth) over other users, and when the packets can come through.

FortiGate firewall components

The FortiGate firewall is comprised of many different features that provides flexibility for the specific needs of your network, both now and as it grows. These features are:

- interfaces (including VLANs)
- zones
- unified threat management (UTM)
- firewall addresses (this includes IPv4 and IPv6, IP pools, wildcard addresses and netmasks, and geography-based addresses)
- monitoring traffic
- traffic shaping and per-ip traffic shaping (advanced)
- firewall schedules
- services (such as AOL, DHCP and FTP)
- logging traffic (advanced)
- QoS (advanced)
- identity-based policies (advanced)
- endpoint security (advanced)

All of these components each provide an important role in configuring your FortiGate firewall. For example, the administrator applies the PING admin access to the wan1 interface so that he or she can ping this external interface and verify that Internet traffic is hitting the internal to wan1 security policy. If there was no PING admin access applied to the external interface, the administrator could not properly verify if traffic is hitting the policy.

For more in-depth explanations of these components, see the [“Firewall components” on page 23](#).

How the firewall components create a FortiGate firewall and help in protecting your network

The firewall components each help in protecting your network, as well as helping traffic to flow better through the network, for example traffic shaping helps to load balance traffic on your network.

The following explains how all of the firewall components get combined to create the FortiGate firewall.

1 In *System > Network > Interface*, create VLAN subinterfaces for each department: sales, marketing and engineering.

These VLAN subinterfaces will be grouped into a zone and the zone will then be applied to a security policy.

2 Create a zone for the VLAN subinterfaces.

3 In *Firewall Objects > Address > Address*, create the IP address ranges that are required: one for sales, one for marketing, and one for engineering.

Each of these ranges corresponds to the departments that have these IP address ranges. For example, sales has 172.16.120.100 - 172.16.120.200.

4 Create a firewall schedule that allows sales and marketing Internet access all day; create another firewall schedule that allows engineering access to the Internet only during their lunch break.

By creating two different firewall schedules, you can block access for one group for a specified time period, and allow another group all day access.

5 Group the firewall schedules together so that you can apply them both to a security policy.

6 Create a virtual IP address that will be used to allow Internet users access to a web server on your DMZ network.

7 In *Policy > Policy > Policy*, create the following:

- a security policy that allows Internet users access to the web server
- a security policy that applies the firewall schedule group for Internet access for the sales, marketing and engineering departments (this applies the zone)
- a deny policy that blocks FTP downloads

8 With all the policies now in the list, arrange them so that the most important policies are first, and least important are last. The list order is:

- deny policy
- security policy that allows Internet users access to the web server
- security policy for sales, engineering and marketing that allows Internet access

Now that all the policies are in the correct order, you need to test that all are working properly.

9 To verify that traffic is hitting the policies, verify that there is a packet count increase occurring in the Count column of each of the policies in the policy list. Troubleshoot any issues using the *diagnose sniffer* and *diagnose debug flow* commands in the CLI.

By testing that traffic is hitting the policies that you just created, you can see whether you need to solve any issues or not. When you use the `diagnose` commands, you can see detailed information about the traffic hitting the policy.

10 Back up the configuration after testing and troubleshooting.

By backing up the changes you made to the configuration, you ensure that a current configuration of this FortiGate firewall configuration is available at any time.

Understanding how a packet travels through the FortiGate unit

Directed by security policies, a FortiGate unit screens network traffic from the IP layer up through the application layer of the TCP/IP stack. The FortiGate firewall plays an important role in how the packet travels through the FortiGate unit out to its destination. The following explains how the packet travels through the FortiGate unit and how the FortiGate firewall plays a role in the life of a packet.

The FortiGate unit performs three types of security inspection:

- stateful inspection, that provides individual packet-based security within a basic session state
- flow-based inspection, that buffers packets and uses pattern matching to identify security threats
- proxy-based inspection, that reconstructs content passing through the FortiGate unit and inspects the content for security threats.

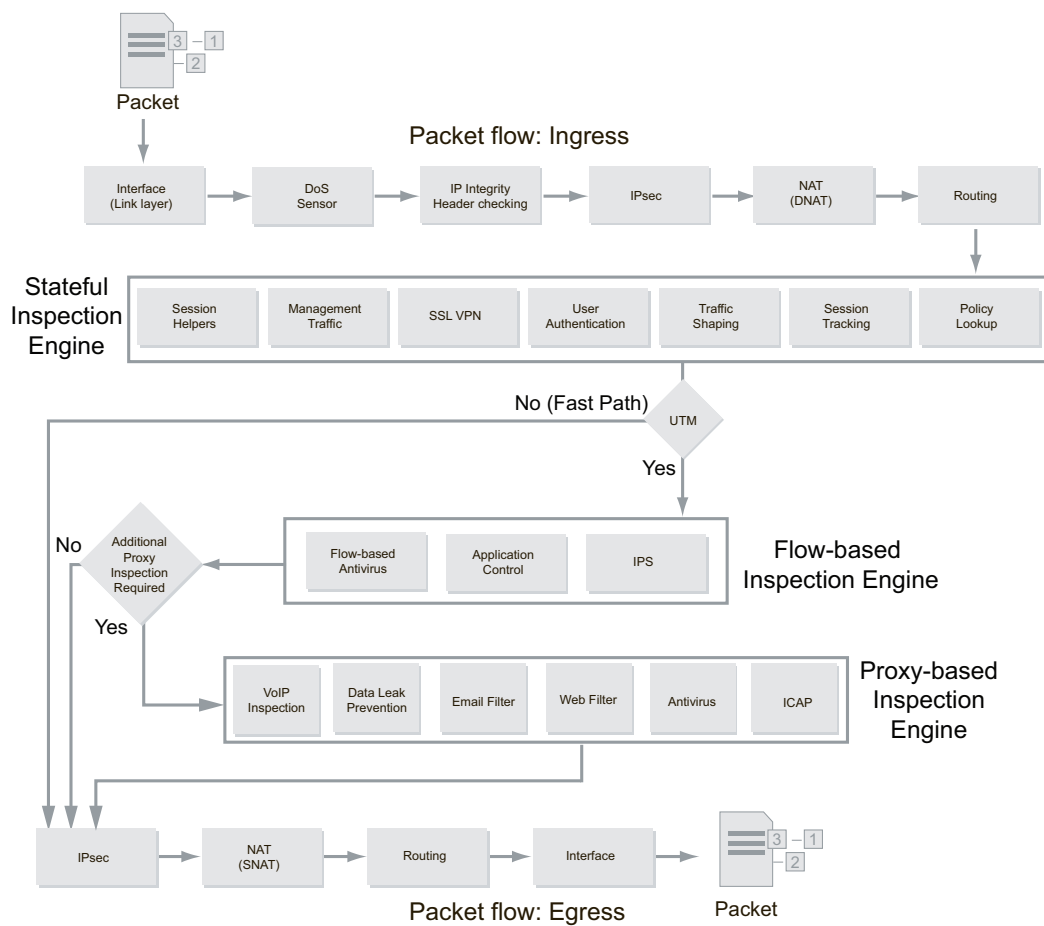
Each inspection component plays a role in the processing of a packet as it traverses the FortiGate unit en route to its destination. When you understand these inspections, you will understand the packet's journey through the FortiGate unit and how the FortiGate firewall helps the packet along to its destination.


For more information about how packets travel through the FortiGate unit, see the Troubleshooting chapter in the [FortiOS Handbook](#). The following explains, in a high-level description, of how a packet travels through the FortiGate unit.

How packets flow in and out of the FortiGate unit

The following provides a high-level description of the steps a packet takes when it enters the FortiGate unit, travelling to its destination, the internal network. Similar steps occur for outbound traffic; they are just in reverse.

- 1 An incoming packet enters the external interface of the FortiGate unit to start its journey through to the internal network. This is called ingress. During ingress, the following processes occur:**
 - DoS Sensor
 - IP integrity header checking
 - IPsec
 - Destination NAT (DNAT)
 - Routing
- 2 After the Routing process finishes, the stateful inspection engine processes the packet, and does the following:**
 - Session Helpers
 - Management Traffic
 - SSL VPN
 - User Authentication
 - Traffic Shaping
 - Session Tracking
 - Policy lookup
- 3 If nothing comes from the stateful inspection engine, then the packet travels to the UTM scanning process. This process may have either a flow-based or proxy-based inspection engine that also processes the packet.**
- 4 If nothing matches the UTM rules, the packet then travels to other processing steps, which include:**
 - IPsec
 - NAT (Source NAT)
 - Routing
 - Internal Interface
- 5 After step 4 is finished, the packet travels out of the internal interface of the FortiGate unit, heading towards its final destination, the internal network. This is referred to as Egress.**

Figure 1: Packet flow



Working with NAT in FortiOS

This section explains NAT and the NAT/Route mode of the FortiGate unit, as well as Transparent mode and its role with NAT. This section also explains the types of NAT that FortiOS supports, including combinations of NAT that you can configure in FortiOS.

This section also includes information about Route mode and how it behaves in FortiOS.

The following topics are included in this section:

- [NAT in FortiOS](#)
- [Types of NAT in FortiOS](#)
- [Combining types of NAT](#)

NAT in FortiOS

Network address translation (NAT) translates one IP address (either a source IP address or destination IP address) for another IP address. NAT in FortiOS, however, can translate IP addresses in many different ways, providing the flexibility you need for your specific network requirements. For example, you can use the Central NAT table to help in translating multiple IP addresses.

When configuring NAT in FortiOS, you should also know how it works within the different modes that the FortiGate unit can be configured in.

This topic contains the following:


- [NAT/Route mode](#)
- [Route mode](#)
- [Transparent mode](#)

NAT/Route mode

In NAT/Route mode, the FortiGate unit is visible to the network that is connected to. All of its interfaces are on different subnets. Each interface it is connected to a network that must be configured with an IP address that is valid for that subnetwork.

NAT/Route mode is typically used when the FortiGate unit is deployed as a gateway between private and public networks. In its default NAT mode configuration, the FortiGate unit functions as a firewall. Security policies control communications through the FortiGate unit to both the Internet and between internal networks. In NAT/Route mode, the FortiGate unit performs network address translation before IP packets are sent to the destination network. For example, a company has a FortiGate unit as their interface to the Internet. The FortiGate unit also acts as a router to multiple sub-networks within the company.

In [Figure 2](#), the FortiGate unit is set to NAT/Route mode and is connected to a network. By using this mode, the FortiGate unit can have a designated port for the Internet, and the internal segments are behind the FortiGate unit, which are invisible to the public access. The FortiGate unit translates IP addresses passing through it to route the traffic to the correct subnet on the Internet.


Figure 2: An example of a FortiGate unit in NAT/Route mode on a network

Route mode

In Route mode, the FortiGate unit is only routing traffic, not translating the IP addresses. In this mode, the FortiGate unit acts similar to a switch, passing the packet along to the destination network. This mode is not to be confused with Transparent mode, which is invisible on the network; rather, in Route mode, the FortiGate unit is visible to the network, but does only routing.

The FortiGate unit is used in Route mode whenever no NAT translation needs to be done. For example, you want to connect two separate subnets without using NAT.

You must select NAT/Route mode when configuring the FortiGate unit for Route mode.

Figure 3: An example of a FortiGate unit in Route mode on a network


Transparent mode

In Transparent mode, the FortiGate unit is invisible to the network. All of its interfaces are on the same subnet and share the same IP address. If you want to configure the FortiGate unit in Transparent mode, all you need to do is to configure a management IP address and a default route.

You would typically use Transparent mode on a private network behind an existing firewall or behind a router. In Transparent mode, the FortiGate unit functions as a firewall and can even perform NAT. Security policies control communications through the FortiGate unit to the Internet and internal network. Traffic cannot pass through until you add security policies when the FortiGate unit is in Transparent mode.

In Transparent mode, you can also perform NAT by creating a security policy or policies that translates the source addresses of packets passing through the FortiGate unit as well as virtual IP addresses and/or IP pools. If you want NAT to be performed in Transparent mode, you must configure two management IP addresses that are on different subnets.

Figure 4: A FortiGate unit in Transparent mode



Types of NAT in FortiOS

There are many types of NAT that are available, some you may already know such as port address translation (PAT). The following explains these types of NAT that are available in FortiOS.

This topic contains the following:

- [Static NAT \(SNAT\)](#)
- [Dynamic NAT \(DNAT\)](#)

Static NAT (SNAT)

Static NAT, or source address translation (SNAT), is when a static source IP address is translated by NAT to another source IP address. In FortiOS, when a packet with a specific source address is accepted by a security policy with NAT enabled, the source address is swapped with another IP address. For example, you want to allow a web server on a private network that is protected by a FortiGate unit to connect to the Internet; the web server has a static IP address of 10.10.30.10 and the external interface of the unit is 172.20.120.233; when the packet is received at the FortiGate's internal interface, it is translated from 10.10.30.10 to 172.20.120.133, and forwards the packet out to the Internet.

Static NAT is used when configuring basic security policies. For example, you want users on a private network to connect to the Internet.

When configuring static NAT security policies, there are several steps that must be configured prior to configuring the actual security policy. For example, for static DNAT, you must configure a virtual IP address that maps to a specific destination address.

Static Destination NAT (SDNAT)

As stated for static NAT, the same is true for static destination address translation, or SDNAT, whereby a packet with a specific destination address is accepted by a security policy with NAT enabled, the destination address is swapped with another destination address.

Static NAT port forwarding

There is also static port forwarding, which acts similarly to static DNAT, translating a destination address and port number to another destination address and port number. The difference is that port forwarding requires a virtual IP address so that the FortiGate unit can properly translate the port number.

When a packet with a destination address to be translated is accepted by a security policy (with DNAT enabled), and a virtual IP with an external port mapped to that address's port, then the FortiGate unit swaps the packet's destination address with the other IP address, and its port number with the external port.

Dynamic NAT (DNAT)

As subnets grow larger, more work is required to set network address translation with each additional client. Rather than assigning static addresses, an administrator may want to set up IP pools. IP pools are ranges of addresses that clients on a subnet can use to send and receive packets, as well as which FortiGate units can use to translate the addresses of packets going through them. This type of translation is known as dynamic NAT, when address translation is done on a flexible or "many-to-one" basis using IP pools.

IP pools do not randomly assign addresses, rather, each IP pool is a prioritized list of IP addresses. When a client is assigned an IP address from the IP pool, it retains that address. Another client that requires an address is then assigned the next IP address from that IP pool list. When the range of virtual IPs are used instead of IP pools, these virtual IPs are prioritized in the same type of list.

Dynamic source address translation

Dynamic source address translation has economies of scale for larger subnets and more flexible subnets, enabling network infrastructure to change without the hassle of reconfiguring addresses after every change. Dynamic source address NAT or DNAT translates many source addresses as defined by an IP pool. Whenever a packet with the specific source address to be translated is accepted by a security policy with source NAT enabled, the FortiGate unit swaps the packet's source address with the other IP address selected from the IP pool.

For example, an organization may want packets leaving the FortiGate unit for the Internet to have source IPs in the range of 172.16.0.1-10. This means that packets accepted by a firewall policy must have their source addresses translated to an address in this range before being forwarded to the Internet. So if the server on the private network with the address 10.0.0.1 has its source IP translated to 172.16.0.1, then the next available source IP in the IP pool will be 172.16.0.2, which a server with address 10.0.0.2 can use.

Dynamic destination address

Dynamic destination address NAT (or DDNAT) translates one range of destination addresses to another range of destination addresses. Whenever a packet within the specified range of destination addresses to be translated is accepted by a security policy with DNAT is enabled, the FortiGate unit swaps the packet's destination address with one of the addresses from the other specified range.

For example, to allow customers from the Internet to connect to several web servers protected by a FortiGate unit, you require a range of Internet addresses (for example, 172.16.0.1-10), enough for each protected web server, and a range of real addresses (for example, 10.0.0.1-10) for each web server. When a packet is received at the external interface of the FortiGate unit with a destination IP address within the Internet range of addresses, the FortiGate unit translates the destination address of the packet to the real address and forwards the packet to the web server on the network protected by the FortiGate unit.

Dynamic port forwarding

Dynamic port forwarding translates one range of destination addresses and ports to another range of destination addresses and ports. Whenever a packet with a specified destination address to be translated is accepted by a security policy with destination NAT enabled, and a virtual IP with an external port mapped to that address's port, then the FortiGate unit swaps the packet's destination address with the other IP address, and its port number with the external port.

For example, to allow customers from the Internet to connect to web servers protected by a FortiGate unit, you require a range of Internet addresses (for example, 172.16.0.1-10) and a range of port numbers (for example, 80-89), and a range of ports numbers to be mapped to (for example, 8080-8089). When a packet is received at the external interface of the FortiGate unit with the 172.16.0.3 destination IP address and port number 8082, the FortiGate unit translates that address to 10.0.0.3 and port number to 82, and then forwards the packet to the web server.

Combining types of NAT

In FortiOS, you can combine a number of NAT features to get the best firewall configuration possible for your network requirements. NAT combinations include Double NAT, which is combining IP pool with virtual IP, and using VIP range for SNAT and static one-to-one mapping.


These combinations can help you when creating your FortiGate firewall configuration. The combinations help when you have multiple addresses (IP pools) and when you need to use a virtual IP address with the IP pool. An example of this combination is called Double NAT.

You can also combine dynamic NAT types, such as dynamic source address translation, to help you with creating the FortiGate firewall using dynamic NAT. An example of this combination is using the Central NAT table.

When considering your FortiGate firewall configuration, you should also consider how to combine NAT types. By combining NAT types, you can easily use multiple addresses when configuring security policies, as well as when you want to provide specific NAT translations, such as using dynamic source NAT that will not change the source port; this combination allows for the handling of specific protocols or services that function only if they use a specific port and that port does not change.

The following are some combinations of NAT that you can use in your FortiGate firewall configuration:

- Double NAT
- Central NAT table (similar to IP pools)
- virtual IP range for SNAT
- static one-to-one mapping
- dynamic source NAT (also known as one-to-one source NAT)
- dynamic source NAT (this uses Dynamic IP pool and a virtual IP)



Firewall components

The FortiGate unit's primary purpose is to act as a firewall to protect your networks from unwanted attacks and to control the flow of network traffic. The firewall consists of many different and important components so that you can better protect your network as your network requirements grow. This section explains these components.

The following topics are included in this section:

- [Using Interfaces and zones in the FortiGate firewall](#)
- [Understanding the firewall address component](#)
- [UTM profiles](#)

Using Interfaces and zones in the FortiGate firewall

Interfaces and zones are used when configuring security policies to define incoming and outgoing traffic. For example, in an internal to wan 1 security policy, the internal interface is where traffic is coming in, and the wan 1 interface is where the traffic is going out to. When the FortiGate unit sees that traffic came in using the internal interface, and needs to leave using the wan 1 (or external interface), the security policy internal to wan1 is matched to the traffic and additional rules are applied to the traffic as well.

Interfaces, either virtual or physical, can be applied to security policies. VLAN subinterfaces are virtual interfaces that can be applied to security policies to control and direct traffic on those subinterfaces. VLAN subinterfaces are interfaces that are part of one of the main interfaces, for example, wan1. For more information about VLAN subinterfaces and how to configure them, see the System Admin chapter of the [FortiOS Handbook](#).

Zones provide the option of grouping multiple FortiGate interfaces, both virtual and physical, that you can then apply to security policies to control the incoming and outgoing traffic on those interfaces. By using zones, you can easily group multiple interfaces and VLAN subinterfaces together to help simplify creating security policies where a number of network segments can use the same policy and UTM settings.

How to apply VLANs and zones and to a security policy

The following explains how to create three VLAN subinterfaces, grouping these subinterfaces into a zone, and then applying the zone to a security policy. The security policy will control the traffic for these VLAN subinterfaces.

1 Create three VLANs in *System > Network > Interface* for engineering, sales and marketing on the internal interface.

These three VLANs will be grouped together to create a zone which will then be applied to the security policy. The zone will be applied to the policy instead of the individual VLANs.

2 Group the VLANs into a zone.

3 Create DHCP servers for each of the VLAN subinterfaces in *System > Network > DHCP*.

4 Create the security policy for the zone to control traffic in *Policy > Policy > Policy*.

In the *Source Interface/Zone* list, you would instead choose the zone. The *Destination Interface/Zone* is the external interface, wan1. By choosing the zone, you apply all the subinterfaces at once.

5 Select *Enable NAT* and *Use Destination Interface Address*; ensure that *Log Allowed Traffic* is also enabled so that you can use the logs to help determine if traffic is hitting the security policy.

Understanding the firewall address component

Firewall addresses in FortiOS provide flexibility when configuring access control over the network traffic. When this document talks about firewall addresses, this encompasses:

- IP addresses and netmasks
- IP pools (this can include the Central NAT table)
- virtual IP addresses
- geography-based addresses
- IPv4 addresses
- wildcard addresses and netmasks
- Fully Qualified Domain Name addresses (FQDN)
- IP address groups

Firewall addresses help define the network addresses that you use when configuring a security policy's source and destination address. The FortiGate unit compares the IP addresses contained in packet headers with a security policy's source and destination addresses to determine if the security policy matches the traffic.

A firewall address can contain one or more network addresses. Network addresses can be represented by an IP address with a netmask, an IP address range, or a fully qualified domain name (FQDN).

When representing hosts by an IP address with a netmask, the IP address can represent one or more hosts. For example, a firewall address can be:

- a single computer, such as 192.45.46.45
- a subnetwork, such as 192.168.1.0 for a class C subnet
- 0.0.0.0, which matches any IP address

The netmask corresponds to the subnet class of the address being added, and can be represented in either dotted decimal or CIDR format. The FortiGate unit automatically converts CIDR formatted netmasks to dotted decimal format. Example formats:

- netmask for a single computer: 255.255.255.255, or /32
- netmask for a class A subnet: 255.0.0.0, or /8
- netmask for a class B subnet: 255.255.0.0, or /16
- netmask for a class C subnet: 255.255.255.0, or /24
- netmask including all IP addresses: 0.0.0.0

Valid IP address and netmask formats include:

- x.x.x.x/x.x.x.x, such as 192.168.1.0/255.255.255.0

- x.x.x.x/x, such as 192.168.1.0/24



An IP address of 0.0.0.0 with a netmask 255.255.255.255 is not a valid firewall address.

When representing hosts by an IP address range, the range indicates hosts with continuous IP addresses in a subnet, such as 192.168.1.[2-10], or 192.168.1.* to indicate the complete range of hosts on that subnet. Valid IP Range formats include:

- x.x.x.x-x.x.x.x, such as 192.168.110.100-192.168.110.120
- x.x.x.[x-x], such as 192.168.110.[100-120]
- x.x.x.*, such as 192.168.110.*

When representing hosts by an FQDN, the domain name can be a subdomain, such as mail.example.com. A single FQDN firewall address may be used to apply a security policy to multiple hosts, as in load balancing and high availability (HA) configurations. FortiGate units automatically resolve and maintain a record of all addresses to which the FQDN resolves. Valid FQDN formats include:

- <host_name>.<second_level_domain_name>.<top_level_domain_name>, such as mail.example.com
- <host_name>.<top_level_domain_name>



Be cautious when employing FQDN firewall addresses. By using a fully qualified domain name in a security policy, while convenient, does present some security risks, because policy matching then relies on a trusted DNS server. If the DNS server should ever be compromised, security policies requiring domain name resolution may no longer function properly.

This topic contains the following:

- [IP addresses for self-originated traffic](#)
- [IP pools](#)
- [IP Pools for security policies that use fixed ports](#)
- [Source IP address and IP pool address matching](#)
- [Geography-based addressing](#)
- [Wildcard addresses](#)
- [Fully Qualified Domain Name addresses](#)
- [Address groups](#)
- [Virtual IP addresses](#)

IP addresses for self-originated traffic

On the FortiGate unit, there are a number of protocols and traffic that is specific to the internal workings of FortiOS. For many of these traffic sources, you can identify a specific port/IP address for this self-originating traffic. The following traffic can be configured to a specific port/IP address:

- SNMP
- Syslog

- alert email
- FortiManager connection IP
- FortiGuard services
- FortiAnalyzer logging
- NTP
- DNS
- Authorization requests such as RADIUS
- FSSO

Configuration of these services is performed in the CLI. In each instance, there is a command `set source-ip`. For example, to set the source IP of NTP to be on the DMZ1 port with an IP of 192.168.4.5, the commands are:

```
config system ntp
  set ntpsyn enable
  set syncinterval 5
  set source-ip 192.168.4.5
end
```

To see which services are configured with source-ip settings, use the `get` command:

```
get system source-ip status
```

The output will appear similar to the sample below:

```
NTP: x.x.x.x
DNS: x.x.x.x
SNMP: x.x.x.x
Central Management: x.x.x.x
FortiGuard Updates (AV/IPS): x.x.x.x
FortiGuard Queries (WebFilter/SpamFilter): x.x.x.x
```

IP pools

An IP pool defines a single IP address or a range of IP addresses. A single IP address in an IP pool becomes a range of one IP address. For example, if you enter an IP pool as 1.1.1.1, the IP pool is actually the address range, 1.1.1.1 to 1.1.1.1. Use IP pools to add NAT policies that translate source addresses to addresses randomly selected from the IP pool, rather than the IP address assigned to that FortiGate interface. You can use the Central NAT table as a way to configure IP pools.

If a FortiGate interface IP address overlaps with one or more IP pool address ranges, the interface responds to ARP requests for all of the IP addresses in the overlapping IP pools.

For example, consider a FortiGate unit with the following IP addresses for the port1 and port2 interfaces:

- port1 IP address: 1.1.1.1/255.255.255.0 (range is 1.1.1.0-1.1.1.255)
- port2 IP address: 2.2.2.2/255.255.255.0 (range is 2.2.2.0-2.2.2.255)

And the following IP pools:

- IP_pool_1: 1.1.1.10-1.1.1.20
- IP_pool_2: 2.2.2.10-2.2.2.20
- IP_pool_3: 2.2.2.30-2.2.2.40

The port1 interface overlap IP range with IP_pool_1 is:

- (1.1.1.0-1.1.1.255) and (1.1.1.10-1.1.1.20) = 1.1.1.10-1.1.1.20

The port2 interface overlap IP range with IP_pool_2 is:

- (2.2.2.0-2.2.2.255) & (2.2.2.10-2.2.2.20) = 2.2.2.10-2.2.2.20

The port2 interface overlap IP range with IP_pool_3 is:

- (2.2.2.0-2.2.2.255) & (2.2.2.30-2.2.2.40) = 2.2.2.30-2.2.2.40

And the result is:

- The port1 interface answers ARP requests for 1.1.1.10-1.1.1.20
- The port2 interface answers ARP requests for 2.2.2.10-2.2.2.20 and for 2.2.2.30-2.2.2.40

Select *Enable NAT* in a security policy and then select *Dynamic IP Pool*. Select an IP pool to translate the source address of packets leaving the FortiGate unit to an address randomly selected from the IP pool.

IP pools cannot be set up for a zone. IP pools are connected to individual interfaces.

IP Pools for security policies that use fixed ports

Some network configurations do not operate correctly if a NAT policy translates the source port of packets used by the connection. NAT translates source ports to keep track of connections for a particular service.

From the CLI you can enable `fixedport` when configuring a security policy for NAT policies to prevent source port translation.

```
config firewall policy
  edit policy_name
    ...
    set fixedport enable
    ...
end
```

However, enabling `fixedport` means that only one connection can be supported through the firewall for this service. To be able to support multiple connections, add an IP pool, and then select *Dynamic IP pool* in the policy. The firewall randomly selects an IP address from the IP pool and assigns it to each connection. In this case, the number of connections that the firewall can support is limited by the number of IP addresses in the IP pool.

Source IP address and IP pool address matching

When the source addresses are translated to the IP pool addresses, one of the following three cases may occur:

Scenario 1: The number of source addresses equals that of IP pool addresses

In this case, the FortiGate unit always matches the IP addressed one to one.

If you enable `fixedport` in such a case, the FortiGate unit preserves the original source port. This may cause conflicts if more than one security policy uses the same IP pool, or the same IP addresses are used in more than one IP pool.

Original address	Change to
192.168.1.1	172.16.30.1
192.168.1.2	172.16.30.2
.....
192.168.1.254	172.16.30.254

Scenario 2: The number of source addresses is more than that of IP pool addresses

In this case, the FortiGate unit translates IP addresses using a wrap-around mechanism.

If you enable `fixedport` in such a case, the FortiGate unit preserves the original source port. But conflicts may occur since users may have different sessions using the same TCP 5 tuples.

Original address	Change to
192.168.1.1	172.16.30.10
192.168.1.2	172.16.30.11
.....
192.168.1.10	172.16.30.19
192.168.1.11	172.16.30.10
192.168.1.12	172.16.30.11
192.168.1.13	172.16.30.12
.....

Scenario 3: The number of source addresses is fewer than that of IP pool addresses

In this case, some of the IP pool addresses are used and the rest of them are not be used.

Original address	Change to
192.168.1.1	172.16.30.10
192.168.1.2	172.16.30.11
192.168.1.3	172.16.30.12
No more source addresses	172.16.30.13 and other addresses are not used

Geography-based addressing

An option is available to add a geography-based address scheme. With this type of addressing, you indicate the geographic region, or country. The FortiGate unit includes an internal list of countries and IP addresses based on historical data from the FortiGuard network.



IPv6 does not support geography-based addressing. This feature is for IPv4 addresses only.

When used in security policies, traffic originating or going to a particular country can be logged, blocked or specific filtering applied.

In the following examples, an geographic-based address for China is added for the WAN1 port.

To add a geography-based address - web-based manager

- 1 Go to *Firewall Objects > Address > Address* and select *Create New*.
- 2 Enter the *Name* of China
- 3 For the *Type*, select *Geography*.

- 4 From the *Country* list, select *China*.
- 5 Select the *Interface* of *WAN1*.
- 6 Select *OK*.

To add a geography-based address - CLI

```
config firewall address
  edit China
    set type geography
    set country CN
    set interface wan1
  end
```

You can use a `diagnose` command to view more information about geography-based addressing. The command displays country and address information for the countries that have been added to firewall addresses.

```
diagnose firewall ipgeo {country-list | ip-list | ip2country}
```

Where:

- `country-list` shows all of the countries that have been added to a firewall address.
- `ip-list` shows the IP addresses of a specified country or all of the countries added to firewall addresses.
- `ip2country` shows the country of origin for a specified IP address. The address must be assigned to one of the countries that has been added to a firewall address.

Wildcard addresses

Wildcard addresses are addresses that identify ranges of IP addresses, reducing the amount of firewall addresses and security policies required to match some of the traffic on your network. Wildcard addresses are an advanced feature, usually required only for complex networks with complex firewall filtering requirements. By using these wildcard addresses in the firewall configuration, administrators can eliminate creating multiple, separate IP addresses and then grouping them to then apply to multiple security policies.

A wildcard address consists of an IP address and a wildcard netmask, for example, 192.168.0.56 255.255.0.255. In this example, the IP address is 192.168.0.56 and the wildcard netmask is 255.255.0.255. The IP address defines the networks to match and the wildcard netmask defines the specific addresses to match on these networks.

In a wildcard netmask, zero means ignore the value of the octet in the IP address, which means the wildcard firewall address matches any number in this address octet. This also means that the number included in this octet of IP address is ignored and can be any number. Usually, if the octet in the wildcard netmask is zero, the corresponding octet in the IP address is also zero.

In a wildcard netmask, a number means match addresses according to how the numbers translate into binary addresses. For example, the wildcard netmask is 255; the wildcard address will only match addresses with the value for this octet that is in the IP address part of the wildcard address. So, if the first octet of the IP address is 192 and the first octet of the wildcard netmask is 255, the wildcard address will only match addresses with 192 in the first octet.

In the above example, the wildcard address 192.168.0.56 255.255.0.255 would match the following IP addresses:

192.168.0.56, 192.168.1.56, 192.168.2.56, ..., 192.168.255.56

The wildcard addresses 192.168.0.56 255.255.0.255 and 192.168.1.56 255.255.0.255 define the same thing since the 0 in the wildcard mask means to match any address in the third octet.

If we use the wildcard address 172.0.20.10 255.0.255.255, it would match the following IP addresses:

172.1.20.10, 172.2.20.10, 172.3.20.10, ..., 172.255.20.10

In a wildcard netmask, a number other than 255 matches multiple addresses for this octet. You can perform a binary conversion to calculate the addresses that would be matched by a given value. For example, to create the IP address and wildcard netmask to match the following network addresses:

```
192.168.32.0/24
192.168.33.0/24
192.168.34.0/24
192.168.35.0/24
192.168.36.0/24
192.168.37.0/24
192.168.38.0/24
192.168.39.0/24
```

Table 1 shows how to write the third octet for these networks according to the octet bit position and address value for each bit.

Table 1: Octet bit position and address value for each bit

Decimal	128	64	32	16	8	4	2	1
32	0	0	1	0	0	0	0	0
33	0	0	1	0	0	0	0	1
34	0	0	1	0	0	0	1	0
35	0	0	1	0	0	0	1	1
36	0	0	1	0	0	1	0	0
37	0	0	1	0	0	1	0	1
38	0	0	1	0	0	1	1	0
39	0	0	1	0	0	1	1	1
	M	M	M	M	M	D	D	D

Since the first five bits match, the networks can be summarized into one network (192.168.32.0/21 or 192.168.32.0 255.255.248.0). All eight possible combinations of the three low-order bits are relevant for the network ranges. The wildcard address that would match all of these subnet addresses can be written as 192.168.32.0 255.255.248.0.

Wildcard addresses are similar to routing access list wildcard masks. You add routing access lists containing wildcard masks using the `config router access-list` command. However, router access list wildcard masks use the inverse of the masking system used for firewall wildcard addresses. For the router access list wildcard masks, zero (0) means match all IP addresses and one (1) means ignore all IP addresses. So to match IP addresses 192.168.0.56, 192.268.1.56, 192.168.2.56, ... 192.168.255.56 you would use the following router access IP address prefix and wildcard mask: 192.168.0.56 0.0.255.0.

Wildcard firewall addresses are configured only in the CLI. The following is an example of how to configure a wildcard firewall address.

```
config firewall address
```

```
edit example_wildcard_address
  set type wildcard
  set wildcard 192.168.0.56 255.255.0.255
end
```

Using wildcard addresses in the firewall configuration

The following example shows how wildcard addresses can be applied to network traffic. This example consists of a security policy where both the source and destination addresses are firewall wildcard addresses.

Source Address: 10.129.5.0 255.127.7.0

Destination Address: 10.129.0.10 255.127.7.255

A security policy with these source and destination addresses would permit:

- A device with IP address 10.129.5.100 to connect through the FortiGate unit to IP address 10.129.0.10
- A device with IP address 10.129.13.100 to connect through the FortiGate unit to IP address 10.129.8.10
- A device with IP address 10.129.21.100 to connect through the FortiGate unit to IP address 10.129.0.10

In another example of wildcard addresses, the following shows how only odd numbered addresses get allowed through:

- 1 Create wildcard address 4.2.2.0/255.255.255.1.

This is configured in the CLI.

- 2 Create a deny security policy that uses the wildcard address, 4.2.2.0.

The results are that only the odd-numbered 4.2.2.0 addresses are allowed in; all other addresses are blocked.

Fully Qualified Domain Name addresses



Be cautious when employing FQDN firewall addresses. Using a fully qualified domain name in a security policy, while convenient, does present some security risks, because policy matching then relies on a trusted DNS server. Should the DNS server be compromised, security policies requiring domain name resolution may no longer function properly.

Using Fully Qualified Domain Name (FQDN) addresses in security policies has the advantage of causing the FortiGate unit to keep track of DNS TTLs and adapt as records change. As long as the FQDN address is used in a security policy, it stores the address in the DNS cache. The FortiGate unit will query the DNS for an amount of time specified, in seconds, and update the cache as required. This feature can reduce maintenance requirements for changing firewall addresses for dynamic IP addresses. This also means that you can create security policies for networks configured with dynamic addresses using DHCP.

You specify the TTL time in the CLI only. For example, to set the TTL for 30 minutes on an FQDN of `www.example.com` on port 1, enter the following commands:

```
config firewall address
  edit FQDN_example
    set type fdqn
    set associated-interface port 1
    set fqdn www.example.com
```

```
set cache-ttl 1800
end
```

Address groups

Similar to zones, if you have a number of addresses or address ranges that require the same security policies, you can put them into address groups, rather than creating multiple similar policies. Because security policies require addresses with homogenous network interfaces, address groups should contain only addresses bound to the same network interface, or to *Any* — addresses whose selected interface is *Any* are bound to a network interface during creation of a security policy, rather than during creation of the firewall address.

For example, if address 1.1.1.1 is associated with port1, and address 2.2.2.2 is associated with port2, they cannot be in the same group. However, if 1.1.1.1 and 2.2.2.2 are configured with an interface of *Any*, they can be grouped, even if the addresses involve different networks.

You cannot mix IPv4 firewall addresses and IPv6 firewall addresses in the same address group.

Virtual IP addresses

In FortiOS, virtual IP addresses (VIPs) can be used when configuring security policies to translate IP addresses and ports of packets received by a network interface. When the FortiGate unit receives inbound packets matching a security policy whose *Destination Address* field is a virtual IP, the FortiGate unit applies NAT, replacing packets's IP addresses with the virtual IP's mapped IP address.

VIPs can specify translation of packets' port numbers and/or IP addresses for both inbound and outbound connections. In Transparent mode, virtual IPs are available only in the CLI.

VIP addresses are typically used to map external (public) to internal (private) IP addresses for Destination NAT (DNAT).

Grouping virtual IPs

You can organize multiple virtual IPs into a virtual IP group to simplify your security policy list. For example, instead of having five identical policies for five different but related virtual IPs located on the same network interface, you might combine the five virtual IPs into a single virtual IP group, which is used by a single security policy.

Security policies using VIP Groups are matched by comparing both the member VIP IP addresses) and port numbers).

Match-vip

The match-vip feature allows the FortiGate unit to log virtual IP traffic that gets implicitly dropped. This feature eliminates the need to create two policies for virtual IPs; one that allows the virtual IP, and the other to get proper log entry for DROP rules.

For example, you have a virtual IP security policy and enabled the match-vip feature; the virtual IP traffic that is not matched by the policy is now caught.

The match-vip feature is available only in the CLI. Use the following command syntax to enable this feature. By default, it is disabled.

```
config firewall policy
```



```
edit <vip_policy_name>
    set match-vip {disable | enable}
end
```

How to use match-vip

In this example, a deny security policy has already been configured that blocks FTP sessions. A virtual IP address will be configured in this example and then applied to a security policy that allows Internet users access to a web server on the company's DMZ network.

1 Create the virtual IP address in *Firewall Objects > Virtual IP > Virtual IP*.

This address is called vip-dmz. You can configure the virtual IP address solely in the CLI. This would eliminate having to go back and forth.

2 Log in to the CLI and enter the following commands:

```
config firewall policy
    edit vip-dmz
        set match-vip enable
    end
```

3 Create the virtual IP security policy.

For this security policy, you need to turn on logging within the security policy.

4 Test the policy to view the activity that is occurring with the `match-vip` command enabled.

Services

Services represent typical traffic types and application packets that pass through the FortiGate unit. Firewall services define one or more protocols and port numbers associated with each service. Security policies use service definitions to match session types. You can organize related services into service groups to simplify your security policy list.

Many well-known traffic types have been predefined in firewall services and protocols on the FortiGate unit. These predefined services and protocols are defaults, and cannot be edited or removed. However, if you require different services, you can create custom services.

To view the predefined servers, go to *Firewall Objects > Service > Predefined*.

If there is a service that does not appear on the list, or you have a unique service or situation, you can create your own custom service. You need to know the ports, IP addresses or protocols of that particular service or application uses, to create the custom service.

Predefined service list

Many well-known traffic types have been predefined in firewall services. These predefined services are defaults, and cannot be edited or removed. However, if you require different services, you can create custom services.

Predefined services are located in *Firewall Objects > Service > Predefined*. [Table 2](#) lists the FortiGate firewall predefined services.

Table 2: Predefined services

Service name	Description	Protocol	Port
AFS3	Advanced File Security Encrypted File, version 3, of the AFS distributed file system protocol.	TCP	7000-7009
		UDP	7000-7009
AH	Authentication Header. AH provides source host authentication and data integrity, but not secrecy. This protocol is used for authentication by IPSec remote gateways set to aggressive mode.	IP	51
ANY	Matches connections using any protocol over IP.	all	all
AOL	America Online Instant Message protocol.	TCP	5190-5194
BGP	Border Gateway Protocol. BGP is an interior/exterior routing protocol.	TCP	179
CVSPSERVER	Concurrent Versions System Proxy Server. CSSP Server is very good for providing anonymous CVS access to a repository.	TCP	2401
		UDP	2401
DCE-RPC	Distributed Computing Environment / Remote Procedure Calls. Applications using DCE-RPC can call procedures from another application without having to know on which host the other application is running.	TCP	135
		UDP	135
DHCP	Dynamic Host Configuration Protocol. DHCP allocates network addresses and delivers configuration parameters from DHCP servers to hosts.	UDP	67 68
DHCP6	Dynamic Host Configuration Protocol for IPv6.	UDP	546, 547
DNS	Domain Name Service. DNS resolves domain names into IP addresses.	TCP	53
		UDP	53
ESP	Encapsulating Security Payload. ESP is used by manual key and AutoIKE IPSec VPN tunnels for communicating encrypted data. AutoIKE VPN tunnels use ESP after establishing the tunnel by IKE.	IP	50
FINGER	A network service providing information about users.	TCP	79
FTP	File Transfer Protocol.	TCP	21
FTP_GET	File Transfer Protocol. FTP GET sessions transfer remote files from an FTP server to an FTP client computer.	TCP	21

Table 2: Predefined services (Continued)

Service name	Description	Protocol	Port
FTP_PUT	File Transfer Protocol. FTP PUT sessions transfer local files from an FTP client to an FTP server.	TCP	21
GOPHER	Gopher organizes and displays Internet server contents as a hierarchically structured list of files.	TCP	70
GRE	Generic Routing Encapsulation. GRE allows an arbitrary network protocol to be transmitted over any other arbitrary network protocol, by encapsulating the packets of the protocol within GRE packets.	IP	47
GTP (FortiOS Carrier only)	GPRS Tunneling protocol (GTP). GTP is used with GSM and UMTS networks to carry user data within GPRS core networks. FortiOS Carrier can accept and process IPv4 GTP packet.	UDP	2123,2152,3386
H323	H.323 multimedia protocol. H.323 is a standard approved by the International Telecommunication Union (ITU) defining how audiovisual conferencing data can be transmitted across networks. For more information, see the FortiGate Support for H.323 Technical Note .	TCP	1720, 1503
		UDP	1719
HTTP	Hypertext Transfer Protocol. HTTP is used to browse web pages on the World Wide Web.	TCP	80
HTTPS	HTTP with secure socket layer (SSL). HTTPS is used for secure communication with web servers.	TCP	443
ICMP_ANY	Internet Control Message Protocol. ICMP allows control messages and error reporting between a host and gateway (Internet).	ICMP	Any
IKE	Internet Key Exchange. IKE obtains authenticated keying material for use with the Internet Security Association and Key Management Protocol (ISAKMP) for IPSEC.	UDP	500, 4500
IMAP	Internet Message Access Protocol. IMAP is used by email clients to retrieve email messages from email servers.	TCP	143
IMAPS	IMAP with SSL. IMAPS is used for secure IMAP communication between email clients and servers. IMAPS is only available on FortiGate units that support SSL content scanning and inspection. For more information, see the UTM chapter of the FortiOS Handbook.	TCP	993
INFO_ADDRESS	ICMP information request messages.	ICMP	17

Table 2: Predefined services (Continued)

Service name	Description	Protocol	Port
INFO_REQUEST	ICMP address mask request messages.	ICMP	15
IRC	Internet Relay Chat. IRC allows users to join chat channels.	TCP	6660-6669
Internet-Locator-Service	Internet Locator Service. ILS includes LDAP, User Locator Service, and LDAP over TLS/SSL.	TCP	389
L2TP	Layer 2 Tunneling Protocol. L2TP is a PPP-based tunnel protocol for remote access.	TCP	1701
		UDP	1701
LDAP	Lightweight Directory Access Protocol. LDAP is used to access information directories.	TCP	389
MGCP	Media Gateway Control Protocol. MGCP is used by call agents and media gateways in distributed Voice over IP (VoIP) systems.	UDP	2427, 2727
MMS (FortiOS Carrier only)	MMS tunneling protocol. MMS is used when sending and receiving multimedia content to a mobile phone.	TCP	1755
		UDP	1024-5000
MS-SQL	Microsoft SQL Server is a relational database management system (RDBMS) produced by Microsoft. Its primary query languages are MS-SQL and T-SQL.	TCP	1433, 1434
MYSQL	MySQL is a relational database management system (RDBMS) which runs as a server providing multi-user access to a number of databases.	TCP	3306
NFS	Network File System. NFS allows network users to mount shared files.	TCP	111, 2049
		UDP	111, 2049
NNTP	Network News Transport Protocol. NNTP is used to post, distribute, and retrieve Usenet messages.	TCP	119
NTP	Network Time Protocol. NTP synchronizes a host's time with a time server.	TCP	123
		UDP	123
NetMeeting	NetMeeting allows users to teleconference using the Internet as the transmission medium.	TCP	1720
ONC-RPC	Open Network Computing Remote Procedure Call. ONC-RPC is a widely deployed remote procedure call system.	TCP	111
		UDP	111
OSPF	Open Shortest Path First. OSPF is a common link state routing protocol.	IP	89
PC-Anywhere	PC-Anywhere is a remote control and file transfer protocol.	TCP	5631
		UDP	5632

Table 2: Predefined services (Continued)

Service name	Description	Protocol	Port
PING	Ping sends ICMP echo request/replies to test connectivity to other hosts.	ICMP	8
PING6	Ping6 sends ICMPv6 echo request/replies to network hosts to test IPv6 connectivity to other hosts.	ICMP6	58
POP3	Post Office Protocol v3. POP retrieves email messages.	TCP	110
POP3S	Post Office Protocol v3 with secure socket layer (SSL). POP3S is used for secure retrieval of email messages. POP3S is only available on FortiGate units that support SSL content scanning and inspection. For more information, see the UTM chapter of the FortiOS Handbook.	TCP	995
PPTP	Point-to-Point Tunneling Protocol. PPTP is used to tunnel connections between private network hosts over the Internet. Note: Also requires IP protocol 47.		47
		TCP	1723
QUAKE	Quake multi-player computer game traffic.	UDP	26000, 27000, 27910, 27960
RADIUS	Remote Authentication Dial In User Service. RADIUS is a networking protocol that provides centralized access, authorization and accounting management for people or computers to connect and use a network service.	TCP	1812, 1813
RAUDIO	RealAudio multimedia traffic.	UDP	7070
RDP	Remote Desktop Protocol is a multi-channel protocol that allows a user to connect to a networked computer.	TCP	3389
REXEC	Rexec traffic allows specified commands to be executed on a remote host running the rexecd service (daemon).	TCP	512
RIP	Routing Information Protocol. RIP is a common distance vector routing protocol. This service matches RIP v1.	UDP	520
RLOGIN	Remote login traffic.	TCP	513
RSH	Remote Shell traffic allows specified commands to be executed on a remote host running the rshd service (daemon).	TCP	514

Table 2: Predefined services (Continued)

Service name	Description	Protocol	Port
RTSP	Real Time Streaming Protocol is a protocol for use in streaming media systems which allows a client to remotely control a streaming media server, issuing VCR-like commands such as play and pause, and allowing time-based access to files on a server.	TCP	554, 7070, 8554
		UDP	554
SAMBA	Server Message Block. SMB allows clients to use file and print shares from enabled hosts. This is primarily used for Microsoft Windows hosts, but may be used with operating systems running the Samba daemon.	TCP	139
SCCP	Skinny Client Control Protocol. SCCP is a Cisco proprietary standard for terminal control for use with voice over IP (VoIP).	TCP	2000
SIP	Session Initiation Protocol. SIP allows audiovisual conferencing data to be transmitted across networks. For more information, see the Voice Solutions: SIP chapter of the FortiOS Handbook.	UDP	5060
SIP-MSNmessenger	Session Initiation Protocol used by Microsoft Messenger to initiate an interactive, possibly multimedia session.	TCP	1863
SMTP	Simple Mail Transfer Protocol. SMTP is used for sending email messages between email clients and email servers, and between email servers.	TCP	25
SMTPS	SMTP with SSL. Used for sending email messages between email clients and email servers, and between email servers securely. SMTPS is only available on FortiGate units that support SSL content scanning and inspection. For more information, see the UTM chapter of the FortiOS Handbook .	TCP	465
SNMP	Simple Network Management Protocol. SNMP can be used to monitor and manage complex networks.	TCP	161-162
		UDP	161-162
SOCKS	SOCKeTS. SOCKS is an Internet protocol that allows client-server applications to transparently use the services of a network firewall.	TCP	1080
		UDP	1080

Table 2: Predefined services (Continued)

Service name	Description	Protocol	Port
SQUID	A proxy server and web cache daemon that has a wide variety of uses that includes speeding up a web server by caching repeated requests; caching web, DNS and other computer network lookups for a group of people sharing network resources; aiding security by filtering traffic.	TCP	3128
SSH	Secure Shell. SSH allows secure remote management and tunneling.	TCP	22
		UDP	22
SYSLOG	Syslog service for remote logging.	UDP	514
TALK	Talk allows conversations between two or more users.	UDP	517-518
TCP	Matches connections using any TCP port.	TCP	0-65535
TELNET	Allows plain text remote management.	TCP	23
TFTP	Trivial File Transfer Protocol. TFTP is similar to FTP, but without security features such as authentication.	UDP	69
TIMESTAMP	ICMP timestamp request messages.	ICMP	13
TRACEROUTE	A computer network tool used to determine the route taken by packets across an IP network.	TCP	33434
		UDP	33434
UDP	Matches connections using any UDP port.	UDP	0-65535
UUCP	Unix to Unix Copy Protocol. UUCP provides simple file copying.	UDP	540
VDOLIVE	VDO Live streaming multimedia traffic.	TCP	7000-7010
VNC	Virtual Network Computing. VNC is a graphical desktop sharing system which uses the RFB protocol to remotely control another computer.	TCP	5900
WAIS	Wide Area Information Server. WAIS is an Internet search protocol which may be used in conjunction with Gopher.	TCP	210
WINFRAME	WinFrame provides communications between computers running Windows NT, or Citrix WinFrame/MetaFrame.	TCP	1494
WINS	Windows Internet Name Service is Microsoft's implementation of NetBIOS Name Service (NBNS), a name server and service for NetBIOS computer names.	TCP	1512
		UDP	1512

Table 2: Predefined services (Continued)

Service name	Description	Protocol	Port
X-WINDOWS	X Window System (also known as X11) can forward the graphical shell from an X Window server to X Window client.	TCP	6000-6063

Service groups

You can organize multiple firewall services into a service group to simplify your security policy list. For example, instead of having five identical policies for five different but related firewall services, you might combine the five services into a single address group that is used by a single security policy.

Service groups can contain both predefined and custom services. Service groups cannot contain other service groups.

You can organize multiple firewall services into a service group to simplify your security policy list. For example, instead of having five identical policies for five different but related firewall services, you might combine the five services into a single service group that is used by a single security policy.

Service groups can contain both predefined and custom services. Service groups cannot contain other service groups.

Firewall schedules

When you add security policies on a FortiGate unit, those policies are always on, policing the traffic through the device. Firewall schedules control when policies are in effect, that is, when they are on. You can create one-time schedules which are schedules that are in effect only once for the period of time specified in the schedule. You can also create recurring schedules that are in effect repeatedly at specified times of specified days of the week.

You can create a recurring schedule that activates a policy during a specified period of time. For example, you might prevent game playing during office hours by creating a recurring schedule that covers office hours.

If a recurring schedule has a stop time that is earlier than the start time, the schedule will take effect at the start time but end at the stop time on the next day. You can use this technique to create recurring schedules that run from one day to the next. For example, to prevent game playing except at lunchtime, you might set the start time for a recurring schedule at 1:00 p.m. and the stop time at 12:00 noon. To create a recurring schedule that runs for 24 hours, set the start and stop times to 00.

You can organize multiple firewall schedules into a schedule group to simplify your security policy list. For example, instead of having five identical policies for five different but related firewall schedules, you might combine the five schedules into a single schedule group that is used by a single security policy.

Schedule groups can contain both recurring and one-time schedules. Schedule groups cannot contain other schedule groups.

Schedule groups

You can organize multiple firewall schedules into a schedule group to simplify your security policy list. For example, instead of having five identical policies for five different but related firewall schedules, you might combine the five schedules into a single schedule group that is used by a single security policy.

Schedule groups can contain both recurring and on-time schedules. Schedule groups cannot contain other schedule groups.

Schedule expiry

The schedule in a security policy enables certain aspects of network traffic to occur for a specific length of time. What it does not do however, is police that time. That is, the policy is active for a given time frame, and as long as the session is open, traffic can continue to flow.

For example, in an office environment, Skype use is allowed between noon and 1pm. During that hour, any Skype traffic continues. As long as that session is open, after the 1pm end time, the Skype conversations can continue, yet new sessions will be blocked. Ideally, the Skype session should close at 1pm.

Using a CLI command you can set the schedule to terminate all sessions when the end time of the schedule is reached. Within the `config firewall` command enter the command:

```
set schedule-timeout enable
```

By default, this is set to disable.

UTM profiles

Where security policies provide the instructions to the FortiGate unit as to what traffic is allowed through the device, the Unified Threat Management (UTM) profiles provide the screening that filters the content coming and going on the network. The UTM profiles enable you to instruct the FortiGate unit what to look for in the traffic that you don't want, or want to monitor, as it passes through the device.

A UTM profile is a group of options and filters that you can apply to one or more firewall policies. UTM profiles can be used by more than one security policy. You can configure sets of UTM profiles for the traffic types handled by a set of security policies that require identical protection levels and types, rather than repeatedly configuring those same UTM profile settings for each individual security policy.

For example, while traffic between trusted and untrusted networks might need strict antivirus protection, traffic between trusted internal addresses might need moderate antivirus protection. To provide the different levels of protection, you might configure two separate protection profiles: one for traffic between trusted networks, and one for traffic between trusted and untrusted networks.

UTM profiles are available for various unwanted traffic and network threats. Each are configured separately and can be used in different groupings as needed. You configure UTM profiles in the *UTM* menu and applied when creating a security policy by selecting the UTM profile type.

For more information about configuring profiles that will be used in a security policy, see the UTM chapter of the [FortiOS Handbook](#).

How to use UTM profiles to monitor and protect your network

In this example, UTM profiles help you in monitoring and protecting your network from viruses, email filtering and web filtering. This example uses the default UTM profiles.

- 1 **Locate the security policy that allows access to the Internet (internal -> wan 1) in *Policy > Policy > Policy*.**

- 2 **On the Edit Policy page, select UTM and then select these options: *Enable Antivirus, Enable Web Filter and Enable Email Filter*.**

The FortiGate unit will apply the antivirus, web filter, and email filter settings to the packet if a match is found.

- 3 **Select *OK*.**

When packets enter the FortiGate unit's internal interface, if a packet matches the internal -> wan 1 policy, the FortiGate unit now scans for viruses and applies any web filtering and email filtering rules if there are matches as well.

- 4 **Go to the *eicar.org* web site and download the eicar test file.**

By downloading the eicar test file, you can determine that the antivirus profile is working properly, as well as to see this activity on the AV Monitor page. When attempting to download the file, a web page appears, stating that you are not permitted to download the file. This indicates that the antivirus profile is working properly.

- 5 **Go to *UTM Profiles > Monitor > AV Monitor* to view the virus activity that just occurred.**


On the page, you should see that the eicar test file was detected by the FortiGate unit; you can select the bar in the chart to see more details. This takes you directly to the FortiGuard Virus Encyclopedia.

- 6 **Go to *UTM Profiles > Monitor > Web Monitor* to view the Internet activity that is occurring on your network.**

On the page, you will see a pie chart that displays all HTTP requests and a bar chart that displays all blocked HTTP requests. If you want to view more detailed information about the blocked requests, hover your mouse over a bar; a tool-tip appears stating how many blocked requests occurred for that item. For example, for Virus, it is one blocked request because you tried to download the eicar test file.

- 7 **Go to *UTM Profiles > Monitor > Email Monitor* to view the email activity that is occurring on your network.**

On the page, you will see both a pie chart and a bar chart, similar to the Web Monitor page. The pie chart displays all the email activity and the bar chart displays all the blocked emails for SMTP, POP3, IMAP, and NNTP.



Security policies

Security policies control all traffic attempting to pass through the FortiGate unit, between FortiGate interfaces, zones, and VLAN subinterfaces.

This section explains what security policies are and how they affect all traffic to and from your network. This section also describes how to configure basic policies which are used as a building block to more complex policies, but they enable you to get the FortiGate unit running on the network quickly.

The following topics are included in this section:

- [Security policy overview](#)
- [Policy order](#)
- [Security policies](#)
- [Creating a basic security policy](#)

Security policy overview

Security policies are instructions the FortiGate unit uses to decide connection acceptance and packet processing for traffic attempting to pass through. When the firewall receives a connection packet, it analyzes the packet's source address, destination address, and service (by port number), and attempts to locate a security policy matching the packet.

Security policies can contain many instructions for the FortiGate unit to follow when it receives matching packets. Some instructions are required, such as whether to drop or accept and process the packets, while other instructions, such as logging and authentication, are optional. The FortiGate unit requires one security policy per traffic flow. For example, network traffic must flow from the internal network to the Internet; a security policy is created (internal interface -> external interface) that allows packets to flow freely from the Internet to the internal network, and from the internal network to the Internet.

Policy instructions may include network address translation (NAT), or port address translation (PAT), or by using virtual IPs or IP pools to translate source and destination IP addresses and port numbers.

Policy instructions may also include UTM profiles, which can specify application-layer inspection and other protocol-specific protection and logging, as well as IPS inspection at the transport layer.

You configure security policies to define which sessions will match the policy and what actions the FortiGate unit will perform with packets from matching sessions.

Sessions are matched to a security policy by considering these features of both the packet and policy:

- Source Interface/Zone
- Source Address
- Destination Interface/Zone
- Destination Address

- Schedule and time of the session's initiation
- Service and the packet's port numbers.

If the initial packet matches the security policy, the FortiGate unit performs the configured Action and any other configured options on all packets in the session.

Packet handling actions can be *ACCEPT*, *DENY*, *IPSEC* or *SSL-VPN*.

- *ACCEPT* policy actions permit communication sessions, and may optionally include other packet processing instructions, such as requiring authentication to use the policy, or specifying one or more UTM profiles to apply features such as virus scanning to packets in the session. An *ACCEPT* policy can also apply interface-mode IPsec VPN traffic if either the selected source or destination interface is an IPsec virtual interface.
- *DENY* policy actions block communication sessions, and you can optionally log the denied traffic. If no security policy matches the traffic, the packets are dropped, therefore it is not required to configure a *DENY* security policy in the last position to block the unauthorized traffic. A *DENY* security policy is needed when it is required to log the denied traffic, also called "violation traffic".
- *IPSEC* and *SSL-VPN* policy actions apply a tunnel mode IPsec VPN or SSL VPN tunnel, respectively, and may optionally apply NAT and allow traffic for one or both directions. If permitted by the firewall encryption policy, a tunnel may be initiated automatically whenever a packet matching the policy arrives on the specified network interface, destined for the local private network.


You need to create security policies based on how the network traffic is going to be flowing through the FortiGate unit. For example, a policy for POP3, where the email server is outside of the internal network, traffic should be from an internal interface to an external interface rather than the other way around. It is typically the user on the network requesting email content from the email server and thus the originator of the open connection is on the internal port, not the external one of the email server. This is also important to remember when view log messages as to where the source and destination of the packets can seem backwards.



If you make any changes to existing policies, those changes take effect immediately.

Security policy list details

The security policy table includes, by default, a number of columns to display information about the policy, for example, source, destination, service, and so on. You can add a number of additional columns to the table to view more information about the policies and what is in their configuration. By going to *Policy > Policy > Policy* and selecting the *Column Settings* link, you can add or remove a number of different columns of information to the policy list, and arrange their placement within the table.

Figure 5: Security policy column selection

Viewing security policies

When viewing security policies in the security policy list, you can view them in either *Section View* or *Global View*. In *Section View*, policies are grouped by how the traffic is directed by interface, for example, internal -> wan1. In *Global View*, policies are listed in one large list with no groupings, referred to as interface pairings.

The FortiGate unit will automatically change the view on the policy list page to *Global View* whenever a policy containing *any* in the *Source interface/zone* or *Destination interface/zone* is created. This occurs because the FortiGate unit understands that this particular policy allows or denies traffic on any FortiGate interface, which breaks the original policy sequence order.

Policies are ordered by fixed policies (ones that contain static interfaces) with each interface pairing (for example, port1 -> port2) and each pairing has their own specific policy order, which does not cause any conflict. However, this interface pairing creates a conflict when a policy containing an ANY interface is created, because the FortiGate unit is now unable to determine which policy set to use and which, in the pair's ordering, should traffic be blocked. The FortiGate unit uses Global View to represent its own understanding of the global policy that was created, using this to help determine the action to take.

Policy order

Each time a FortiGate unit receives a connection attempting to pass through one of its interfaces, the unit searches its security policy list for a matching security policy.

The search begins at the top of the policy list and progresses in order towards the bottom. The FortiGate unit evaluates each policy in the security policy list for a match until a match is found. When the FortiGate unit finds the first matching policy, it applies the matching policy's specified actions to the packet, and disregards subsequent security policies. Matching security policies are determined by comparing the security policy and the packet's:

- source and destination interfaces
- source and destination firewall addresses

- services
- time/schedule.

If no policy matches, the connection is dropped.

As a general rule, you should order the security policy list from most specific to most general because of the order in which policies are evaluated for a match, and because only the **first** matching security policy is applied to a connection. Subsequent possible matches are not considered or applied. Ordering policies from most specific to most general prevents policies that match a wide range of traffic from superseding and effectively masking policies that match exceptions.

















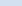
For example, you might have a general policy that allows all connections from the internal network to the Internet, but want to make an exception that blocks FTP. In this case, you would add a policy that denies FTP connections above the general policy.

Figure 6: Example: Blocking FTP — Correct policy order

<input type="checkbox"/>	Status	ID	Source	Destination	Schedule	Service	Action	
Internal -> wan1 (2)								
<input type="checkbox"/>	<input checked="" type="checkbox"/>	2	all	all	always	FTP		}Exception }General
<input type="checkbox"/>	<input checked="" type="checkbox"/>	1	all	all	always	ANY		
Implicit (1)								
<input type="checkbox"/>	Implicit		all	all	always	ANY		

FTP connections would immediately match the deny policy, blocking the connection. Other kinds of services do not match the FTP policy, and so policy evaluation would continue until reaching the matching general policy. This policy order has the intended effect. But if you reversed the order of the two policies, positioning the general policy before the policy to block FTP, all connections, including FTP, would immediately match the general policy, and the policy to block FTP would never be applied. This policy order would not have the intended effect.

Figure 7: Example: Blocking FTP — Incorrect policy order

	Status	ID	Source	Destination	Schedule	Service	Action	
internal -> wan1 (2)								
		1	 all	 all	 always	 ANY		}General }Exception
		2	 all	 all	 always	 FTP		
Implicit (1)								
	Implicit		all	all	always	ANY		

Similarly, if specific traffic requires authentication, IPsec VPN, or SSL VPN, you would position those policies above other potential matches in the policy list. Otherwise, the other matching policies would always take precedence, and the required authentication, IPsec VPN, or SSL VPN might never occur.



A default security policy may exist, which accepts all connections. You can move, disable or delete it. If you move the default policy to the bottom of the security policy list and no other policy matches the packet, the connection will be accepted. If you disable or delete the default policy and no other policy matches the packet, the connection will be dropped.

You can arrange the security policy list to influence the order in which policies are evaluated for matches with incoming traffic. When more than one policy has been defined for the same interface pair, the first matching security policy will be applied to the traffic session.

How to arrange policies

In this example, there are four policies that the FortiGate unit must use when packets enter the FortiGate unit's interface. These policies are IPsec VPN, DENY, Internet access, and an identity-based policy. You need to make sure that the policies are arranged so that the policies that are important do not get left out.

1 On the policy list, select *Global* view to view all policies in the list.

By viewing the list using Global view, you can easily see all policies in the list regardless of the sections that they are in. This helps you to see where in the list you need to move the policies, without having to expand each section to view the policies.

2 Move the IPsec VPN policy to the first line in the table.

You want the IPsec VPN policy to come first so that the process matches this policy first. If the IPsec VPN policy is not first, other policies would always take precedence and the authentication required for IPsec may never occur.

3 Move the DENY policy to the third line of the table.

This DENY policy contains information that denies all FTP traffic.

4 Move the identity-based policy to the fourth line in the table.

5 Move the Internet access policy after the identity-based policy.

Security policies

There are many different security policies that you can configure for the FortiGate firewall. These policies include SSL VPN, wireless, and identity-based policies. With different configurations come different security policies, and each contain different information for processing the packets coming into the FortiGate unit.

The following explain each type of security policy that can be configured and the reason for configuring such a security policy.

This topic contains the following:

- [Identity-based policies](#)
- [SSL VPN policies](#)
- [IPsec policies](#)
- [Accept policies](#)
- [Deny policies](#)
- [IPv6 policies](#)
- [Security policy 0](#)
- [Local-in policies](#)



If you make any changes to existing policies, those changes take effect immediately.

Identity-based policies

If you enable *Enable Identity Based Policy* in a security policy, network users must send traffic involving a supported firewall authentication protocol to trigger the firewall authentication challenge, and successfully authenticate, before the FortiGate unit will allow any other traffic matching the security policy.

User authentication can occur through any of the following supported protocols:

- HTTP
- HTTPS
- FTP
- Telnet

Authentication can also occur through automatic login using NTLM and FSSO receiverships, to bypass user intervention.

The authentication style depends on which of these supported protocols you have included in the selected firewall services group and which of those enabled protocols the network user applies to trigger the authentication challenge. The authentication style will be one of two types. For certificate-based (HTTPS or HTTP redirected to HTTPS only) authentication, you must install customized certificates on the FortiGate unit and on the browsers of network users, which the FortiGate unit matches. For user name and password-based (HTTP, FTP, and Telnet) authentication, the FortiGate unit prompts network users to input their firewall user name and password.

For example, if you want to require HTTPS certificate-based authentication before allowing SMTP and POP3 traffic, you must select a firewall service (in the security policy) that includes SMTP, POP3 and HTTPS services. Prior to using either POP3 or SMTP, the network user would send traffic using the HTTPS service, which the FortiGate unit would use to verify the network user's certificate; upon successful certificate-based authentication, the network user would then be able to access his or her email.

In most cases, you should ensure that users can use DNS through the FortiGate unit without authentication. If DNS is not available, users will not be able to use a domain name when using a supported authentication protocol to trigger the FortiGate unit's authentication challenge.



If you do not install certificates on the network user's web browser, then network users may see an SSL certificate warning message and have to manually accept the default FortiGate certificate, which the network user's web browser may then deem as invalid.



When you use certificate authentication, if you do not specify any certificate when you create a security policy, the FortiGate unit will use the default certificate from the global settings. If you specify a certificate, the per-policy setting will override the global setting.

Authentication requires that *Action* is ACCEPT or SSL-VPN, and that you first create users, assign them to a firewall user group, and assign UTM profiles to that user group.

Identity-based policy example

With this basic identity-based policy example, the security policy will allow HTTPS traffic passing from the external interface (WAN1) to the internal interface (Internal) at all times, as soon as the network user enters their user name and password. For simplicity, the policy will request the firewall authentication. This authentication can be set up for users by going to *User > User > User* and their groupings by going to *User > User Group > User Group*. For this example, the group “accounting” is used. When a user attempts to browse to a secure site, they will be prompted for their log in credentials.

To create a identity-based policy - web-based manager

- 1 Go to *Policy > Policy > Policy* and select *Create New*.
- 2 Enter the following:
- 3 Select *Enable Identity Based Policy*.
- 4 *Firewall* authentication is enabled by default.
- 5 Select *Add*.
- 6 From the *Available User Groups* list, select the *Accounting* user group and select the right arrow to move it to the *Selected User Groups* area.
- 7 From the *Available Services* list, select the *HTTPS* and select the right arrow to move it to the *Selected Services* area.
- 8 For the *Schedule*, select *Always*.
- 9 Select *OK*.

To create a identity-based policy - CLI

```
config firewall policy
  edit 1
    set srcintf internal
    set srcaddr 10.13.20.22
    set dstintf wan1
    set dstaddr 172.20.120.141
    set action accept
    set schedule always
    set identity-based enable
    config identity-based-policy
      edit 1
        set group accounting
        set service HTTPS
        set schedule always
      end
    end
  end
```

SSL VPN policies

SSL VPN security policies are created for permitting SSL VPN clients, web-mode or tunnel-mode, access to the protected network behind the FortiGate unit. These security policies also contain authentication information that will authenticate the users and user group or groups.

IPsec policies

IPsec policies allow IPsec VPN traffic access to the internal network from a remote location. These policies include authentication information that authenticates users and user group or groups. These policies specify the following:

- the FortiGate interface that provides the physical connection to the remote VPN gateway, usually an interface connected to the Internet
- the FortiGate interface that connects to the private network
- IP addresses associated with data that has to be encrypted and decrypted
- optional: a schedule that restricts when the VPN can operate, and services (or types of data) that can be sent.



For a route-based (interface mode) VPN, you do not configure an IPsec security policy. Instead, you configure two regular ACCEPT security policies, one for each direction of communication, with the IPsec virtual interface as the source or destination interface, as appropriate.

Accept policies

Accept security policies accept traffic that is coming into the network. These policies allow traffic through the FortiGate unit, where the packets are scanned, translated if NAT is enabled, and then sent out to its destination.

Accept security policies are the most common security policies that are created in FortiOS. These security policies are basic policies, such as allowing Internet access, as well as complex policies, such as IPsec VPN.

For information about how to configure accept policies, see [“Security policy list details” on page 46](#).

Deny policies

Deny security policies deny traffic that is coming into the network. The FortiGate unit automatically blocks traffic that is associated with a deny security policy.

Deny security policies are usually configured when you need to restrict specific traffic, for example, SSH traffic. Deny security policies can also help when you want to block a service, such as DNS, but allow a specific DNS server.

For information about how to configure DENY policies, see [“Security policy list details” on page 46](#).

How to allow DNS queries to only one DNS server

In this example, a specific DNS server is used for all DNS queries. All other requests for DNS is not allowed. A deny security policy is used to restrict this access.

1 In *Firewall Objects > Address > Address*, create an IP address for the DNS server.

This address will be used for the policy that allows DNS requests from this DNS server.

2 Create a new security policy that blocks all DNS sessions to the Internet.

This policy would have the Action set to DENY and the Service set to DNS. In this policy, the FortiGate unit restricts all requests for any DNS queries.

3 Create a new policy that allows access to only the DNS server.

This policy is used by the FortiGate unit to allow DNS requests to the DNS server that is specified.

4 Move the policies so that they are in the correct order.

If the policies are not in the correct order, the FortiGate unit will not process the instructions properly and the policies will not work properly. The allowed policy needs to be first and the deny policy needs to come right after.

5 Test the policies.

You can test the policies by using diagnose debug command in the CLI or view the packet count in the Count columns of the policies. For more information about how to test and/or verify if traffic is hitting a policy, see [“How to create a basic security policy for Internet access” on page 55.](#)

IPv6 policies

IPv6 security policies are created both for an IPv6 network, and a transitional network. A transitional network is a network that is transitioning over to IPv6, but must still have access to the Internet or must connect over an IPv4 network.

These policies allow for this specific type of traffic to travel between the IPv6 and IPv4 networks. The IPv6 options for creating these policies is hidden by default. You must enable this feature in *System > Admin > Settings*.

For more information about IPv6 in FortiOS, see [“Internet Protocol version 6 \(IPv6\)” on page 63.](#)

Security policy 0

Any security policy that is automatically added by the FortiGate unit has a policy ID number of zero (0). The most common reasons the FortiGate unit creates this policy is:

- The IPsec policy for FortiAnalyzer (and FortiManager version 3.0) is automatically added when an IPsec connection to the FortiAnalyzer unit or FortiManager is enabled.
- The policy to allow FortiGuard servers to be automatically added has a policy ID number of zero.
- The (default) drop rule that is the last rule in the policy and that is automatically added has a policy ID number of zero.
- When a network zone is defined within a VDOM, the intra-zone traffic set to allow or block is managed by policy 0 if it is not processed by a configured security policy.

This policy can appear in logs but will never appear in the security policy list, and therefore, can never be repositioned in the list.

When viewing the FortiGate logs, you may find a log field entry indicating policyid=0. The following log message example indicates the log field policyid=0 in bold.

```
2008-10-06 00:13:49 log_id=0022013001 type=traffic
subtype=violation pri=warning vd=root SN=179089 duration=0
user=N/A group=N/A rule=0 policyid=0 proto=17 service=137/udp
app_type=N/A status=deny src=10.181.77.73 srcname=10.181.77.73
dst=10.128.1.161 dstname=10.128.1.161 src_int=N/A
dst_int="Internal" sent=0 rcvd=0 src_port=137 dst_port=137 vpn=N/A
tran_ip=0.0.0.0 tran_port=0
```

Local-in policies

Security policies control the flow of traffic through the FortiGate unit. The FortiGate unit also includes the option of controlling internal traffic, that is, management traffic.

Each interface includes an allow access configuration to allow management access for specific protocols. Local policies are set up automatically to allow all users all access. Local-in policies takes this a step further, to enable or restrict the user with that access. This also extends beyond the allow access selection.

Local-in policies are configured in the CLI with the commands:

```
config firewall local-in-policy
edit <policy_number>
set intf <source_interface>
set srcaddr <source_address>
set dstaddr <destination_address>
set action {accept | deny}
set service <service name>
set schedule <schedule_name>
end
```

For example, you can configure a local-in policy so that only administrators can access the FortiGate unit on weekends from a specific management computer at 192.168.21.12 using SSH on port 3 (192.168.21.77) using the Weekend schedule which defines the time of access.

```
config firewall local-in-policy
edit <1>
set intf port3
set srcaddr 192.168.21.12
set dstaddr 192.168.21.77
set action accept
set service SSH
set schedule Weekend
end
```

You can also disable a policy should there be a requirement to turn off a policy for troubleshooting or other purpose. To disable a policy enter the commands:

```
config firewall local-in-policy
edit <policy_number>
set status disable
end
```

Use the same commands with a status of `enable` to use the policy again.

Local-in policies are also supported for IPv6 by entering the command `config firewall local-in-policy6`.

Creating a basic security policy

The following describes how to configure a basic security policy as well as how to test and verify that traffic hitting the policy.

This topic includes the following:

- [How to create a basic security policy for Internet access](#)
- [How to verify if traffic is hitting the basic security policy](#)
- [How to test the basic security policy](#)

How to create a basic security policy for Internet access

The following explains how a basic security policy is created, as well as how to test and verify that the policy is working properly. Testing a policy and verifying if traffic is hitting a policy are two ways to ensure that the policy that you created is working properly.

1 In the web-based manager, go to *Policy > Policy > Policy* and select *Create New*.

2 The source interface should be *internal* and the destination interface should *wan1*.

This indicates to the FortiGate unit that the incoming packets will be coming from the internal network and proceeding to the public network or Internet. The interfaces are also understood in reverse: packets that are coming from the outside or Internet and are destined for the internal network.

3 The source and destination addresses should *all*.

This is the default IP address range in *Firewall Objects > Addresses > Address*. This default IP address range indicates that any IP address is accepted within the range. This is written as 0.0.0.0/0.0.0.0.

4 For this policy, you must choose the default *always* schedule for *Schedule*, the *ANY* service for *Service*, and the *Action* to *ACCEPT*.

The default schedule always provides the time limitation, which is none, for the policy. A time limitation can limit the access users have to the Internet or can allow users to access resources at any time of the day or night.

5 Select *Log Allowed Traffic* to view the traffic activity using either *Policy > Monitor > Policy Monitor*, or traffic logs. Select *OK* to save the security policy.

You should test the policy after it has been created. To test a security policy, go to a web site; if you are able to get to the web site, the policy is working properly. You can also view the Count column on the Policy page. The Count column displays the number of packets that have recently passed through, which increases as the packets pass through the FortiGate unit.

How to test the basic security policy

After a security policy has been configured, you can test to see if the policy is working. This should be done after you create a security policy so that you can modify the policy's settings, if required, before backing up the configuration. You should always back up the configuration after making modifications to the FortiGate configuration; by doing so, you will have a current configuration whenever you need it.

1 On a computer that is on the internal network, open a web browser and access any web site.

You should be able to get to that web site.

2 If you are unable to get to a web site, use the following to help troubleshoot the problem:

- Is the policy order correct?
- Using the diag debug flow command, see if traffic is hitting the policy. If not, use the diag sniffer command to determine what is going on
- View the Count column; if no number appears, traffic is not hitting the policy.

3 After troubleshooting the problem, browse to a web site and if you can access it, and then save the current configuration.

How to verify if traffic is hitting the basic security policy

After configuring a security policy, you will want to verify that it is working properly. The following explains how to verify that traffic is hitting the basic security policy that you configured in [“How to verify if traffic is hitting the basic security policy” on page 56](#).

1 In the web-based manager, go to *Policy > Policy > Policy* and locate the internal to wan1 policy.

2 In the Count column, verify that there is packets hitting the security policy.

The Count column displays the amount of packets that are hitting the security policy. In the beginning this count will be low, be increase as the packets come through the FortiGate unit.

3 Go to *Policy > Monitor > Policy Monitor* to view the security policy.


On the Policy Monitor page, you can see the active sessions, bytes or packets that are occurring from the bar chart and table. By selecting the bar within the chart, you can view more detailed information.

4 Go to the CLI, log in, and use `diag debug flow` commands to show traffic is hitting the security policy.

The `diag debug flow` commands show packet flow through the FortiGate unit. The following is an example of what the information gives when you use the `diag debug flow` commands to see if traffic is hitting a policy.

```
diagnose debug enable
diagnose debug flow show console enable
diagnose debug flow filter add 192.168.1.110
diagnose debug flow trace start 50

id=36871 trace_id=1 msg="vd-root received a packet(proto=6,
    192.168.1.110:3152->172.16.100.148:80) from internal."
id=36871 trace_id=1 msg="allocate a new session-0000724b"
id=36871 trace_id=1 msg="find a route: gw-172.20.120.2 via wan1"
id=36871 trace_id=1 msg="find SNAT: IP-172.20.120.11, port-
    40156"
id=36871 trace_id=1 msg="Allowed by Policy-3: SNAT"
id=36871 trace_id=1 msg="SNAT 192.168.1.110-
    >172.20.120.11:40156"
```

Monitoring firewall traffic

You can easily monitor the network traffic on the FortiGate firewall from either the Dashboard or the Monitor menus in the Policy and Firewall Objects menus. By using these monitors, you can understand how to improve your firewall, or resolve issues.

The following explains the various features that you can use to monitor firewall traffic.

The following topics are included in this section:

- [Session tables](#)
- [Monitoring security policy traffic activity](#)

Session tables

Firewall session tables include entries to record source and destination IP addresses and port numbers. For each packet received by a FortiGate unit, it references the session table for a match. Packets of an established session are checked against the session table continually throughout the communication. The performance of depends on the performance of processing session table.

Firewall sessions clear from the table based on the timeout, that is, Time-to-live (TTL) setting. Equally, a completely inactive session with no FIN or RESET will be flushed by the by the session TTL timer. Sessions are not closed based on FIN or a RESET. A FIN that is acknowledged with a FIN ACK would slush the session.

Viewing session tables in the web-based manager

Firewall sessions are viewable in the web-based manager using the *Top Sessions* widget. If this widget is not on the Dashboard, select the *Widget* link at the top of the web-based manager and select it from the pop-up dialog box.

While this view shows a graph of the connecting users and IP addresses, double-clicking on a bar in the graph will display the complete session information for that user.

You can clear a session from the table by scrolling to the right and selecting the delete icon for a given session.

Sessions Monitor

Session information display in *Policy > Monitor > Session Monitor*. You can delete sessions, refresh so that you are viewing current sessions, and you can also filter the session information on the page as well. Filtering allows you to view specific information. For example, you want to view only TCP sessions.

Session Monitor page

Displays the sessions that are currently being monitored by the unit.

Refresh Select to refresh the information in the list.

Filter Settings	<p>Select to filter the information on the page. <i>Filters</i> appears automatically after selecting <i>Filter Settings</i>, below the column headings. Use to configure filter settings.</p> <p>Note: <i>Filter Settings</i> configures all filter settings. Filter icons are used to configure filter settings within that column.</p> <p>To apply a filter setting, select the plus sign beside <i>Add new filter</i> and then select and enter the information required. Repeat to add other filter settings.</p> <p>To modify the settings, select <i>Change</i> beside the setting and edit the settings.</p> <p>To clear all filters settings. Select the icon beside <i>Clear all filters</i>.</p> <p>To use a filter icon to filter settings within a column, select the filter icon in the column. <i>Filters</i> appears. Within <i>Filters</i>, configure the settings for that column.</p>
IPv4	Select to display only IPv4 addresses.
IPv6	Select to display only IPv6 addresses.
Both	Select to display both IPv4 and IPv6 addresses.
Total Concurrent Sessions: <number>/ New Sessions per Second: <number>	Indicates the total number of concurrent sessions, as well as new sessions that are occurring each second.
Page Controls	Use to navigate through the list.
Total: <number>	The total number of current sessions.
#	The number of the session within the list.
Protocol	The service protocol of the connection, for example, UDP.
Src Address	The source IP address of the connection.
Src Port	The source port of the connection.
Src NAT IP	The source NAT IP address.
Src NAT Port	The source NAT IP port.
Dst Address	The destination address of the connection.
Dst Port	The destination port of the connection.
Policy ID	The security policy identification number.
Expiry (sec)	The time, in seconds, before the connection expires.
Duration (sec)	The duration, in seconds, of the session.
Delete	Select to remove a session from within the list.

Viewing session tables in the CLI

Session tables and information is also viewable from the CLI. More information on sessions are available from the CLI where various diagnose commands reveal more granular data. To view the session information enter the following CLI command:

```
diagnose sys session list
```

Output will look something similar to:

```

session info: proto=17 proto_state=01 duration=121 expire=58
              timeout=0 flags=00000000 sockflag=00000000 sockport=0
              av_idx=0 use=4
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0 hakey=0
policy_dir=0 tunnel=/
state=may_dirty br
statistic(bytes/packets/allow_err): org=63/1/1 reply=133/1/1
              tuples=2
origin->sink: org pre->post, reply pre->post dev=6->2/2->6
              gwy=0.0.0.0/0.0.0.0
hook=pre dir=org act=noop 172.20.120.85:51167-
              >8.8.8.8:53(0.0.0.0:0)
hook=post dir=reply act=noop 8.8.8.8:53-
              >172.20.120.85:51167(0.0.0.0:0)
misc=0 policy_id=3 id_policy_id=0 auth_info=0 chk_client_info=0
              vd=0
serial=000171db tos=ff/ff app_list=0 app=0
dd_type=0 dd_rule_id=0
per_ip_bandwidth meter: addr=172.20.120.85, bps=1984
total session 189

```

To clear a session enter the following command:

```
diagnose sys session clear
```

State	Meaning
log	Session is being logged
local	Session is originating from, or destined for, a local stack.
ext	Session is created by a firewall session helper.
may_dirty	Session is created by a policy. For example, the session for FTP channel control will have this state but the FTP data channel will not.
ndr	Session will be checked by an IPS signature.
nds	Session will be checked by an IPS anomaly.
br	Session is being bridged, that is, in transparent mode.
npu	Session will possibly be offloaded to NPU.
wccp	Session is handled by WCCP.

Proto_state fields: TCP

The proto_state field value has two digits. This is because the FortiGate unit keeps track of the original direction and the reply direction.

State	Value	Expire Timer Default (seconds)
NONE	0	10
ESTABLISHED	1	3600
SYN_SENT	2	120

SYN & SYN/ACK	3	60
FIN_WAIT	4	120
TIME_WAIT	5	120
CLOSE	6	10
CLOSE_WAIT	7	120
LAST_ACK	8	30
LISTEN	9	120

Proto_state fields: SCTP

State	Value	Expire Timer Default (seconds)
SCTP_S_NONE	0	60
SCTP_S_ESTABLISHED	1	3600
SCTP_S_CLOSED	2	10
SCTP_S_COOKIE_WAIT	3	5
SCTP_S_COOKIE_ECHOED	4	10
SCTP_S_SHUTDOWN_SENT	5	30
SCTP_S_SHUTDOWN_REC'D	6	30
SCTP_S_ACK_SENT	7	3
SCTP_S_MAX	8	120

Proto_state fields: UDP

UDP is a sessionless protocol, however the FortiGate unit still monitors two different states:

- Reply Not Seen - 0
- Reply Seen - 1

In the example output below, a state of 00, the UDP packet has been seen and a session will be created, but no reply packet has been seen:

```
session info: proto=17 proto_state=00 expire=179 timeout=3600
use=3
```

In this example, the UDP packet has been seen and a session created. Reply packets have also been seen:

```
session info: proto=17 proto_state=01 expire=22 timeout=3600
use=3
```

Proto_state field for ICMP

There are no states for ICMP traffic; it will always appear as `proto_state=00`.

Monitoring security policy traffic activity

The Policy Monitor page provides information about the activity of security policies. This activity can be viewed at a high level, or in much more detail, by drilling down to get more specific information.

The Policy Monitor page allows you to view the information in either a graphical format, or in a table. The graphical format, or chart, provides an easy and user-friendly view of the traffic activity that is occurring. The chart also provides a way to drill-down to more information; you can view this information by selecting on a bar within the chart. The drill-down information can be displayed by source address or destination address or by destination port.

Below the chart, a table provides information as well about each policy include the type of action the policy

Policy Monitor page


Displays information about the security policy traffic occurring on the unit.

Tip: View additional and more detailed information by selecting a bar within the chart.

Refresh	Select to refresh the information on the page.
Reset	Select to reset the information to clear the current information from the page. New information is included on the page.
Top Policy Usage	Displays the top security policy usage in a bar chart.
Report By	Select to view information by the current active sessions, bytes or packets.

(Table explaining detailed information about the top policy usage)

Policy ID	The security policy identification number.
Source Interface/Zone	The source address or zone used within that security policy.
Destination Interface/Zone	The destination address of zone used within that security policy.
Action	The type of action that is specified in the security policy. For example <i>Action</i> is set to <i>DENY</i> . The action displays as an icon; for example, a green check mark is <i>ALLOW</i> .
Bytes	The number of bytes used by the security policy. This is reflected in the bar chart.
Packets	The number of packets.



Internet Protocol version 6 (IPv6)

This section explains IPv6 in FortiOS. This section does not explain IPv6 in its entirety, only a high-level summary of IPv6 and how IPv6 is supported in FortiOS. For any additional information about IPv6, see the ipv6.com web site.

The following topics are included in this section:

- [What is IPv6?](#)
- [IPv6 in FortiOS](#)
- [Dual stack routing configuration](#)
- [IPv4 tunneling configuration](#)
- [Remotely connecting to an IPv6 network over the Internet](#)
- [IPv6 overview](#)
- [Transition from IPv4 to IPv6](#)
- [Configuring FortiOS to connect to an IPv6 tunnel provider](#)
- [FortiGate IPv6 configuration](#)
- [IPv6 troubleshooting](#)
- [FortiGate IPv6 configuration](#)
- [IPv6 troubleshooting](#)
- [Additional IPv6 resources](#)

What is IPv6?

Internet Protocol version 6 (IPv6) is the next-generation version of IP addressing. This updated version of IP addressing provides many advances, such as more routing efficiency and reducing the need for NAT. IPv6 also provides better security and mobility support, as well as stateless auto-reconfiguration of hosts which allows IPv6 hosts to automatically configure when connected to a routed IPv6 network.

IPv6 uses 128-bit addressing, which is written in hexadecimal digits separated by a colon. For example, 2001:DB8::6334. This revised version of IP addressing has the potential to provide trillions and trillions of addresses, or an address for each device on the Internet.



For IPv6 address examples, documents use the IPv6 special address 2001:DB8::/32 to indicate that the address is an example. This is stated in RFC 3849. For more information about the specific addresses that are used in IPv6, see ipv6.com.

IPv6 in FortiOS

By default, the FortiGate unit is not enabled to use IPv6 options and settings; however, they are there. To enable IPv6, go to *System > Admin > Settings* and select *IPv6 Support on GUI*. When enabled, you can use IPv6 addressing on any of the address-dependant components of the FortiGate unit, including security policies, interface addressing and DNS servers. IPv6 addressing can be configured on the web-based manager and in the CLI.

There are many different features that FortiOS supports in IPv6. The following is what FortiOS supports in IPv6:

- Static routing
- Dynamic routing (RIPv6, BGP4+, and OSPFv3)
- DNS
- Network interface addressing
- Routing access lists and prefix lists
- IPv6 tunnel over IPv4 and IPv4 tunnel over IPv6
- Security policies
- Authentication
- IPv6 over SCTP
- UTM protection
- Packet and network sniffing
- IPsec VPN
- SSL VPN
- UTM protection
- NAT/Route and Transparent mode
- Logging and reporting
- SNMP
- Virtual IPs and groups
- IPv6-specific troubleshooting, such as ping6

When configuring IPv6 in FortiOS, you can create a dual stack route or IPv4-IPv6 tunnel. A dual stack routing configuration implements dual IP layers, supporting both IPv4 and IPv6, in both hosts and routers. An IPv4-IPv6 tunnel is essentially similar, creating a tunnel that encapsulates IPv6 packets within IPv4 headers that carry these IPv6 packets over IPv4 tunnels. The FortiGate unit can also be easily integrated into an IPv6 network.

IPv6 works almost the same as IPv4 in FortiOS. The only main difference is the IP addresses, since you are using IPv6 addressing instead of IPv4. There is also no NAT, unless you are configuring a dual stack routing or IPv4 tunnelling configuration.

Connecting the FortiGate unit to an IPv6 network is exactly the same as connecting it to an IPv4 network, the only difference is that you are using IPv6 addresses.

Dual stack routing configuration

A dual stack routing configuration implements dual IP layers in hosts and routers, supporting both IPv6 and IPv4. The FortiOS dual stack architecture supports both IPv4 and IPv6 traffic and routes the appropriate traffic as required to any device on the network. Administrators can update network components and applications to IPv6 on their own schedule, and even maintain some IPv4 support indefinitely if that is necessary.

Devices that are on this type of configured network, and connect to the Internet, can query Internet DNS servers for both IPv4 and IPv6 addresses. If the Internet site supports IPv6, the device can easily connect using the IPv6 address. If the Internet site does not support IPv6, then the device can connect using the IPv4 addresses. The dual stack architecture of FortiOS provides all the features that you need for protecting your network, such as UTM security for the traffic, and routing.

If an organization with a mixed network uses an Internet service provider that does not support IPv6, they can use an IPv6 tunnel broker to connect to IPv6 addresses that are on the Internet. FortiOS supports IPv6 tunneling over IPv4 networks to tunnel brokers. The tunnel broker extracts the IPv6 packets from the tunnel and routes them to their destinations.

IPv4 tunneling configuration

In an IPv4 tunneling configuration, IPv6 packets are encapsulated within IPv4 headers, which carry these IPv6 packets over IPv4 tunnels. This type of configuration is more appropriate for those who have completely transitional over to IPv6, but need an Internet connection, which is still mostly IPv4 addresses.

Remotely connecting to an IPv6 network over the Internet

Similar to the IPv4 tunneling configuration, FortiOS supports IPv6 tunneling over IPv4 across the Internet between two IPv6 networks that are protected by FortiGate units.

All traffic between the IPv6 networks are tunnelled over IPv4, which in this case is the Internet. Each FortiGate unit extracts the IPv6 traffic from the IPv4 tunnel and traffic on the internal networks uses IPv6.

In FortiOS, you configure this type of network configuration using IPsec VPN because IPv6 is supported for IPsec VPNs. The VPN provides higher security for the data transmitted between the IPv6 networks. This configuration includes an interface-based IPsec VPN between IPv6 interfaces on each FortiGate unit.

IPv6 overview

IP version 6 handles issues that weren't around decades ago when IPv4 was created such as running out of IP addresses, fair distributing of IP addresses, built-in quality of service (QoS) features, better multimedia support, and improved handling of fragmentation. A bigger address space, bigger default packet size, and more optional header extensions provide these features with flexibility to customize them to any needs.

IPv6 has 128-bit addresses compared to IPv4's 32-bit addresses, effectively eliminating address exhaustion. This new very large address space will likely reduce the need for network address translation (NAT) since IPv6 provides more than a billion IP addresses for each person on Earth. All hardware and software network components must support this new address size, an upgrade that may take a while to complete and will force IPv6 and IPv4 to work side-by-side during the transition period. During that time FortiOS supports IPv4 and IPv6 will ensure a smooth transition for networks.

Differences between IPv4 and IPv6

Table 3: IPv4 and IPv6 differences

Property	IPv4	IPv6
Address size	32 bits	128 bits
Network size	8 - 30 bits	64 bits
Packet header size	20 - 60 bytes	40 bytes
Header-level extension	Limited number of small IP options.	Unlimited number of IPv6 extension headers.

Table 3: IPv4 and IPv6 differences

Property	IPv4	IPv6
Fragmentation	Sender or any intermediate router allowed to fragment.	Only sender may fragment.
Control Protocols	Mixture of non-IP (ARP), ICMP and other protocols.	All control protocols based on ICMPv6.
Minimum MTU	567 bytes	1280 bytes
Address assignment	one address per host	multiple addresses per interface.
Address types	Use of unicast, multicast and broadcast address types.	Broadcast addressing no longer used, use of unicast, multicast and anycast address types
Address configuration	Devices configured manually or with host configuration protocols such as DHCP.	Devices configure themselves independently using stateless auto configuration or use DHCP.

IPv6 addresses are assigned to interfaces rather than nodes, thereby recognizing that a node can have more than one interface, and you can assign more than one IPv6 address to an interface. In addition, the larger address space in IPv6 addresses allows flexibility in allocating addresses and routing traffic, and simplifies some aspects of address assignment and renumbering when changing Internet Service Providers.

With IPv4, complex Classless Inter-Domain Routing (CIDR) techniques were developed to make the best use of the small address space. CIDR facilitates routing by allowing blocks of addresses to be grouped together into a single routing table entry. With IPv4, renumbering an existing network for a new connectivity provider with different routing prefixes is a major effort (see [RFC 2071, Network Renumbering Overview: Why would I want it and what is it anyway?](#) and [RFC 2072, Router Renumbering Guide](#)). With IPv6, however, it is possible to renumber an entire network ad hoc by changing the prefix in a few routers, as the host identifiers are decoupled from the subnet identifiers and the network provider's routing prefix.

The size of each subnet in IPv6 is 2^{64} addresses (64 bits), which is the square of the size of the entire IPv4 Internet. The actual address space utilized by IPv6 applications will most likely be small in IPv6, but both network management and routing will be more efficient.

IPv6 MTU

Maximum Transmission Unit (MTU) refers to the size (in bytes) of the largest packet or frame that a given layer of a communications protocol can pass onwards. A higher MTU brings higher bandwidth efficiency. IPv6 requires an MTU of at least 1280 bytes. With encapsulations (for example, tunneling), an MTU of 1500 or more is recommended.

IPv6 address format

The IPv6 address is 128 bits long and consists of eight, 16-bit fields. Each field is separated by a colon and must contain a hexadecimal number. In [Figure 8](#), an X represents each field.

The IPv6 address is made up of two logical parts:

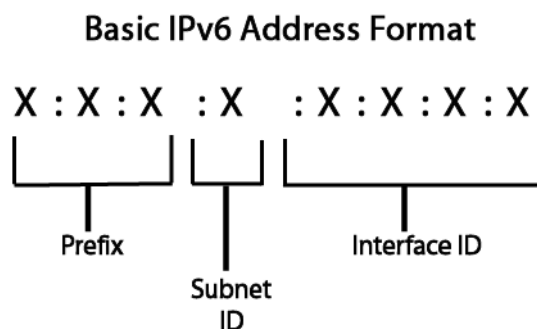
- 64-bit (sub)network prefix

- 64-bit host

The (sub)network prefix part contains the site prefix (first three fields, 48 bits) and the subnet ID (next two fields, 16-bits), for a total of 64-bits. The information contained in these fields is used for routing IPv6 packets. The (sub)network prefix defines the site topology to a router by specifying the specific link to which the subnet has been assigned. The site prefix details the public topology allocated (usually by an Internet Service Provider, ISP) to your site. The subnet ID details the private topology (or site topology) to a router that you assign to your site when you configure your IPv6 network.

The host part consists of the interface ID (or token) which is 64-bits in length and must be unique within the subnet. The length of the interface ID allows for the mapping of existing 48-bit MAC addresses currently used by many local area network (LAN) technologies such as Ethernet, and the mapping of 64-bit MAC addresses of IEEE 1394 (FireWire) and other future LAN technologies. The host is either configured automatically from the MAC address of the interface, or is manually configured.

Figure 8: IPv6 Address Format



IP address notation

IPv6 addresses are normally written as eight groups of four hexadecimal digits each, separated by a colon, for example:

2001:db8:3c4d:0d82:1725:6a2f:0370:6234

is a valid IPv6 address.

There are several ways to shorten the presentation of an IPv6 address. Most IPv6 addresses do not occupy all of the possible 128 bits. This results in fields that are “padded” with zeros or contain only zeros. If a 4-digit group is 0000, it may be replaced with two colons (::), for example:

2001:db8:3c4d:0000:1725:6a2f:0370:6234

is the same IPv6 address as:

2001:db8:3c4d::1725:6a2f:0370:6234

Leading zeroes in a group may be omitted, for example (in the address above):

2001:db8:3c4d::1725:6a2f:370:6234

The double colon (::) must only be used once in an IP address, as multiple occurrences lead to ambiguity in the address translation.

The following examples of shortened IP address presentations all resolve to the same address.

19a4:0478:0000:0000:0000:0000:1a57:ac9e
 19a4:0478:0000:0000:0000::1a57:ac9e
 19a4:478:0:0:0:0:1a57:ac9e

```
19a4:478:0:0::1a57:ac9e
19a4:478::0:0:1a57:ac9e
19a4:478::1a57:ac9e
```

All of these address presentations are valid and represent the same address.

For IPv4-compatible or IPv4-mapped IPv6 addresses (see [“Address types” on page 68](#)), you can enter the IPv4 portion using either hexadecimal or dotted decimal, but the FortiGate CLI always shows the IPv4 portion in dotted decimal format. For all other IPv6 addresses, the CLI accepts and displays only hexadecimal.

Netmasks

As with IP addresses, hexadecimal notation replaces the dotted decimal notation of IPv4. IPv4 Classless Inter-Domain Routing (CIDR) notation can also be used. This notation appends a slash (“/”) to the IP address, followed by the number of bits in the network portion of the address.

Table 4: IPv6 address notation

IP Address	3ffe:ffff:1011:f101:0210:a4ff:fee3:9566
Netmask	ffff:ffff:ffff:ffff:0000:0000:0000:0000
Network	3ffe:ffff:1011:f101:0000:0000:0000:0000
CIDR IP/Netmask	3ffe:ffff:1011:f101:0210:a4ff:fee3:9566/64

Address scopes

Address scopes define the region where an address may be defined as a unique identifier of an interface. The regions are: local link (link-local), site network (site-local), and global network. Each IPv6 address can only belong to one zone that corresponds to its scope.

Address types

IPv6 addresses are classified into three groups - [Unicast](#), [Multicast](#), and [Anycast](#).

Unicast

Identifies an interface of an individual node. Packets sent to a unicast address are sent to that specific interface. Unicast IPv6 addresses can have a scope reflected in more specific address names - global unicast address, link-local address, and unique local unicast address.

Multicast

Multicast addresses are assigned to a group of interfaces that typically belong to different nodes. A packet that is sent to a multicast address is delivered to all interfaces identified by the address.

IPv6 multicast addresses are distinguished from unicast addresses by the value of the high-order octet of the addresses. A value of 0xFF (binary 11111111) identifies an address as a multicast address. Any other value identifies an address as a unicast address.

The four least significant bits of the second address octet identify the address scope or the span over which the multicast address is propagated.

Anycast

Anycast addresses are assigned to a group of interfaces usually belonging to different nodes. A packet sent to an anycast address is delivered to just one of the member interfaces, typically the 'nearest' according to the router protocols' choice of distance. They cannot be identified easily as their structure is the same as a normal unicast address, differ only by being injected into the routing protocol at multiple points in the network. When a unicast address is assigned to more than one interface (making it an anycast address), the address assigned to the nodes must be configured in such a way as to indicate that it is an anycast address.

Interfaces configured for IPv6 must have at least one link-local unicast address and additional ones for site-local or global addressing. Link-local addresses are often used in network address autoconfiguration where no external source of network addressing information is available.

Special addresses

Special IPv6 addresses include unspecified and loopback addresses. For more information about IPv6 addresses, see [RFC 4921, IP Version 6 Addressing Architecture](#)

The IPv6 address space is split into scopes, or address scopes. The table below indicates which IPv6 address is used.

Table 5: IPv6 addresses with prefix information

Address Type	Binary Prefix	IPv6 Notation	Uses
Embedded IPv4 address	00...1111 1111 1111 1111 (96 bits)	::FFF/96	Prefix for embedding IPv4 address in an IPv6 address.
Loopback	00...1 (128 bits)	::1/128	Used as a node to send an IPv6 packet to itself. Seen as link-local unicast address of a virtual interface (loopback interface) to an imaginary link that goes nowhere. Must never be assigned to a physical interface, or as the source address of IPv6 packets that are sent outside of the single node. IPv6 destination address of loopback should not be sent outside a single node, and never forwarded by an IPv6 router. Equivalent to 127.0.0.1 in IPv4. RFC 246022
Global unicast	001	2000::3	Global unicast and anycast. RFC 429120
Global unicast	01 - 1111 1000 0	4000::/2 - FC00::/9	Global unicast and anycast (unallocated)

Table 5: IPv6 addresses with prefix information

Address Type	Binary Prefix	IPv6 Notation	Uses
Teredo	0010 0000 000 0001 0000 0000 000 0000	2001:0000::/32	Teredo - RFC 438023
Nonroutable	0010 0000 0000 0001 1101 1000 1000 0000	2001:D88::/32	Nonroutable. Documentation purposes only - RFC 384924
6to4	0010 0000 0000 0010	2002::/16	Used for communication between two nodes running both IPv4 and IPv6 over the Internet. Formed by combining the IPv6 prefix with the 32-bits of the public IPv4 address of the node, creating a 48-bit address prefix. - RFC 3056
6Bone	0011 1111 1111 1110	3FFE::/16	Deprecated. 6Bone testing assignment 1996 to mid-2006 RFC 370125
Local-link unicast	1111 1110 10	FE80::/10	Used for addressing on a single link for automatic address configuration, neighbor discovery, or when no routers are present. Routers must not forward packets with link-local source or destination addresses.
Reserved	1111 1110 11	FEC0::/10	Used for addressing inside of a site without needing a global prefix. Routers must not forward packets with site-local source or destination addresses outside of the site. RFC 387926
Local IPv6 address	1111 110	FC00::/7	Unicast unique local address space, unicast and anycast - RFC 419327
Multicast	1111 1111	FF00::/8	Multicast address space - RFC 4291 For more information, see "Multicast" on page 68 .

Header Extension

The base header of an IPv6 address is fixed for efficient processing. Header extensions are indicated by the next header value in the next header field. Header extensions are optional and do not need to be present in all IPv6 packets. The sequence for the next header in order is represented by the diagram below.

IPv6 Header	Hop-by-Hop Options Header	Destinations Options Header Router	Routing Header	Fragment Header	Authentication on Header	Encapsulation Security Payload	Destination Options Header Destination	Mobility Header (MIPv6)	TCP/UDP/ Sctp	Payload
-------------	---------------------------	------------------------------------	----------------	-----------------	--------------------------	--------------------------------	--	-------------------------	---------------	---------

The last header extension is the value of either 6 for TCP, 17 for UDP, 132 for SCTP or any other transport protocol defined by the IETF.

Header extensions appear in the following sequence:

- Hop-by-Hop Options Header
 - First Extension Header
 - Next Header value of 0 indicates the Hop-by-Hop Options Extension Header
 - All nodes along the route or path must process this extension header
- Routing Header
 - Second Extension Header
 - Next Header value of 43 indicates the Routing Extension Header
 - All nodes along the route or path must process this extension header
 - Note that the Routing Header Type 0 is due to security reason depreciated
- Fragmentation Header
 - Third Extension Header
 - Next Header value of 44 indicates the Fragmentation Extension Header
 - Used in case of transmitting payload longer than a IPv6 packet can carry
- Authentication Header (AH)
 - Fourth Extension Header
 - Next Header value of 50 indicates the Authentication Extension Header
 - Used to provide protection against replay, origin authentication and connectionless integrity
- Encapsulating Security Payload (ESP) Header
 - Fifth Extension Header
 - Next Header value of 51 indicates the ESP Extension Header
 - Used to provide protection against replay, origin authentication and connectionless integrity
- Destination Options Header
 - Sixth Extension Header
 - Next Header value of 60 indicates the Destination Options Header
 - Used to provide additional information for the end systems node
- Mobility Header
 - Seventh Extension Header
 - Next Header value of 135 indicates the Mobility Header
 - Used by mobile nodes to exchange information for Mobile IP nodes (MIPv6)

Table 6: Header and Protocol Types

Extension Header	Type
------------------	------

Table 6: Header and Protocol Types

Hop-by hop Options	0
Routing	43
Fragment	44
Destination Options	60
Authentication Header (AH)	50
Encapsulating Security Payload	51
Mobility	135
Protocol	Type
TCP	6
UDP	17
IPv6-in-IPv6	41
GRE	47
ICMPv6	58
No next header	59
OSPF	89
PIM	103
SCTP	132

IPv6 neighbor discovery

IPv6 Neighbor Discovery (ND) is a set of messages and processes that determine relationships between neighboring nodes. Neighboring nodes are on the same link. The IPv6 ND protocol replaces the IPv4 protocols Address Resolution Protocol (ARP), Internet Control Message Protocol (ICMPv4), Router Discovery (RDISC), and ICMP Redirect, and provides additional functionality. The IPv6 ND protocol facilitates the autoconfiguration of IPv6 addresses. Autoconfiguration is the ability of an IPv6 host to automatically generate its own IPv6 address, making address administration easier and less time-consuming.

Hosts use ND to:

- discover addresses, address prefixes, and other configuration parameters
- discover neighboring routers.

Routers use ND to:

- advertise their presence, host configuration parameters, and on-link prefixes
- inform hosts of 'better' next-hop address to forward packets for a specified destination.

Nodes use ND to:

- resolve link-layer address of a neighboring node to which an IPv6 packet is being forwarded and determine whether the link-layer address of a neighboring node has altered
- determine whether IPv6 packets can be sent to and received from a neighbor
- automatically configure IPv6 addresses for its interfaces.

To facilitate neighbor discovery, routers periodically send messages advertising their availability. This communication includes lists of the address prefixes for destinations available on each router's interfaces.

ND defines five different Internet Control Message Protocol (ICMP) packet types: a pair of Neighbor Solicitation and Neighbor Advertisement messages, a pair of Router Solicitation and Router Advertisement messages, and a Redirect message.

A Neighbor Solicitation is sent by a node to determine the link-layer address of a neighbor, or to verify that a neighbor is still reachable via a cached link-layer address. Also used for Duplicate Address Detection (how a node determines that an address it wants to use is not already in use by another node). The Neighbor Advertisement message is a response to a Neighbor Solicitation message. A node may also announce a link-layer address change by sending unsolicited Neighbor Advertisements.

A host may send a Router Solicitation when an interface becomes enabled, requesting routers to generate a Router Advertisement immediately rather than at their next scheduled time.

Routers advertise their presence together with various link and Internet parameters according to a specific schedule or in response to a Router Solicitation message. A Router Advertisement contains prefixes used for on-link determination and/or address configuration, a suggested hop limit value, etc.

The Redirect message is used by routers to inform hosts of a better first-hop for a destination.

For more information, see RFC 2461, *Neighbor Discovery for IP Version 6 (IPv6)*.

Transition from IPv4 to IPv6

If the Internet is to take full advantage of the benefits of IPv6, there must be a period of transition to enable IPv6-only hosts to reach IPv4 services and to allow isolated IPv6 hosts and networks to reach the IPv6 Internet over the IPv4 infrastructure.

RFC 2893, *Transition Mechanisms for IPv6 Hosts and Routers* and RFC 2185, *Routing Aspects of IPv6 Transition* define several mechanisms to ensure that IPv6 hosts and routers maintain interoperability with the existing IPv4 infrastructure, and facilitate a gradual transition that does not impact the functionality of the Internet. The mechanisms, known collectively as Simple Internet Transition (SIT), include:

- dual-stack IP implementations for hosts and routers that must interoperate between IPv4 and IPv6
- embedding of IPv4 addresses in IPv6 addresses. IPv6 hosts are assigned addresses that are interoperable with IPv4, and IPv4 host addresses are mapped to IPv6
- IPv6-over-IPv4 tunneling mechanisms to encapsulate IPv6 packets within IPv4 headers to carry them over IPv4 infrastructure
- IPv4/IPv6 header translation, used when implementation of IPv6 is well-advanced and few IPv4 systems remain.

FortiGate units are dual IP layer IPv6/IPv4 nodes and they support IPv6 over IPv4 tunneling. For more information, see [RFC 2893, Transition Mechanisms for IPv6 Hosts and Routers](#) and [RFC 2185, Routing Aspects of IPv6 Transition](#).


Configuring FortiOS to connect to an IPv6 tunnel provider

If an organization with a mixed network uses an Internet service provider that does not support IPv6, they can use an IPv6 tunnel broker to connect to IPv6 addresses on the Internet. FortiOS supports IPv6 tunnelling over service provider IPv4 networks to tunnel brokers. The tunnel broker extracts the IPv6 packets from the tunnel and routes them to their IPv6 destination. The internal network is running IPv6. The FortiGate unit creates an IPv6-over-IPv4 tunnel to the IPv6 tunnel broker. From the tunnel broker, your network can access IPv6 addresses on the Internet.

In this example the internal network is small and directly connected to the FortiGate unit. There is no need for routing on the internal network since everything is connected and on the same subnet. For this example, consider the following:

- Before configuring your FortiGate unit for IPv6-over-IPv4 tunneling, you need to choose an IPv6 tunnel broker and get their information.
- The addresses used in this example are for example use only.
- VDOMs are not enabled.
- The tunnel broker IPv4 address is 78.35.24.124.
- The tunnel broker IPv6 end of the tunnel is 2001:4dd0:ff00:15e::1/64.
- The FortiGate unit external IPv4 address is 172.20.120.17.
- The FortiGate unit IPv6 address of the tunnel is 2001:4dd0:ff00:15e::2/64.
- port1 of the FortiGate unit is connected to the internal network.
- port2 of the FortiGate unit is connected to the external network (Internet).

Figure 9: Connecting to an IPv6 tunnel broker



Steps to connect to an IPv6 tunnel broker

- 1 Create a SIT-Tunnel Interface.
- 2 Create a static IPv6 Route into the Tunnel-Interface.

- 3 Assign your IPv6 Network to your FortiGate.
- 4 Create a Firewall-Policy to allow Traffic from LAN to the Tunnel-Interface.

Create a SIT-tunnel interface

Creating the SIT-tunnel creates a virtual interface in the form of a tunnel, much like a VPN interface. The end points of the tunnel are the FortiGate unit and the tunnel broker's server addresses.

In the example, the external address of the FortiGate unit is DHCP-based and may change to any value on that subnet, so the source address allows for that.

```
config system sit-tunnel
  edit HE_ip6_broker
    set destination 78.35.24.124
    set interface port2
    set ip6 2001:4dd0:ff00:15e::2/64
    set source 172.20.120.0
  next
end
```

Now that the tunnel exists, some additional interface commands are required. Such as enabling ping6 for troubleshooting and allow HTTPS and SSH administration connections to the interface.

```
config system interface
  edit HE_ip6_broker
    config ipv6
      set ip6-allowaccess ping https ssh
    end
  next
end
```

Create a static IPv6 route into the tunnel-Interface

With the tunnel up and the security policies in place, all that remains is to add a default route for IPv6 traffic to go over the tunnel. As there will only be one static routing entry, there is no need for a priority. This may change in the future if other routes are added.

```
config router static6
  edit 1
    set device HE_ip6_broker
  next
end
```

Assign your IPv6 network to your FortiGate

This step assigns an IPv6 address to the internal interface on the FortiGate unit. That way all IPv6 traffic entering on this interface will be routed to the tunnel. Systems with addresses within this prefix are reachable on the subnet in question without help from a router, so the `onlink-flag` is enabled. Hosts can create an address for themselves by combining this prefix with an interface identifier, so the `autonomous-flag` is enabled.

```
config system interface
  edit port1
    config ipv6
      set ip6-address 2001:4dd0:ff42:72::1/64
      set ip6-allowaccess ping https ssh
      config ip6-prefix-list
        edit 2001:4dd0:ff42:72::/64
```

```

        set autonomous-flag enable
        set onlink-flag enable
        set preferred-life-time 3600
        set ip6-send-adv enable
    next
end
next
end

```

At this point any PCs on your internal network that are set to auto-configure, should have their addresses. To test this you can ping6 from the PC to the FortiGate unit. See [“IPv6 ping description” on page 88](#).

Create a security policy to allow traffic from port1 to the tunnel interface

With the tunnel configured, it will appear as an interface in the Network interface list. That means the next step is to add a security policies to allow traffic to and from the tunnel.

```

config firewall policy6
edit 2
    set srcintf port1
    set dstintf HE_ip6_broker
    set srcaddr "::/0"
    set dstaddr "::/0"
    set action accept
    set schedule "always"
    set service "ANY"
    set logtraffic enable
next
end

```

Test the connection

To test the tunnel, try to connect to an external IPv6 address such as <http://ipv6.google.com>.

If you want to see the path the IPv6 traffic takes, do a traceroute from a PC on the internal network to an external address. You will see the traffic enter the FortiGate unit, enter the tunnel, pass through the tunnel broker server, and on out over the Internet.

If you are entering an IPv6 address into your web browser, you have to type: [https://\[2001:4dd0:ff42:72::1\]](https://[2001:4dd0:ff42:72::1]). The square brackets are to discriminate between the address part and a port, like in [https://\[2001:4dd0:ff42:72::1\]:8080](https://[2001:4dd0:ff42:72::1]:8080)

FortiGate IPv6 configuration

FortiOS (4.0 MR2) supports the following FortiOS IPv6 features (all configurable from the web-based manager or CLI):

- Static routing and dynamic routing
- Network interface addressing
- DHCP Server (CLI only)
- Routing access lists and prefix lists
- IPv6 tunnel over IPv4, IPv4 tunnel over IPv6
- Security policies and identity-based security policies

- Local-in security policies
- IPv6 over SCTP
- Packet and network sniffing
- IPsec VPNs
- UTM protection including
- NAT/Route and Transparent mode
- Logging and reporting
- IPv6 specific troubleshooting such as ping6

Displaying IPv6 options on the web-based manager

Before configuring IPv6 using the web-based manager, you must first turn on IPv6 display by going to *System > Admin > Settings* and selecting the IPv6 support option. Once turned on, IPv6-related options and pages appear throughout the web-based manager. For example, you can add IPv6 addresses to any FortiGate interface, you can add IPv6 DNS server IP addresses, IPv6 security policies, IPv6 firewall addresses and so on.

UTM protection for IPv6 networks

FortiOS uses IPv6 security policies to provide UTM protection for IPv6 traffic. Antivirus, web filtering, FortiGuard Web Filtering, email filtering, FortiGuard Email Filtering, data leak prevention (DLP), and VoIP protection features can be enabled in IPv6 security policies using normal FortiOS UTM profiles for each UTM feature.

Configuring IPv6 interfaces

The dual stack architecture is most obvious when configuring IPv6 on interfaces on your FortiGate unit.

IPv6 interfaces - web-based manager

In the *Addressing mode* section of the *Create New* or *Edit* screen, there are two fields instead of one. Without IPv6 enabled, there is only the *IP/Netmask* field for IPv4 addresses. With IPv6 enabled, there is an additional field called *IPv6 Address*.

With both addresses configured for an interface, that interface will accept both IPv4 and IPv6 traffic. Each protocol will be handled differently, depending on the security policies and routing in place for it. This allows traffic from IPv6 to be sent to other IPv6 devices, and IPv4 traffic to be sent only to other IPv4 devices. This separation of the traffic is required because if IPv6 traffic is sent to devices that don't support it, that traffic will not reach its destination.

You should enable IPv6 Administrative Access to connect to the IPv6 address of an interface for administration.

IPv6 interfaces - CLI

In the CLI, there are a number of IPv6 specific interface settings. These are found as part of the `config system interface` command under `config ipv6`. In the CLI there are many more settings available, although many are optional. The settings that are required or recommended are highlighted.

```
config system interface
edit <interface_string>
config ipv6
```

```

set ip6-address <ipv6_addr>
set ip6-allowaccess <http https ping ssh telnet>
set ip6-link-mtu <bytes_int>
set ip6-send-adv <enable | disable>
set autoconf <enable | disable>
set ip6-default-life <seconds_int>
set ip6-hop-limit <count_int>
set ip6-manage-flag <enable | disable>
set ip6-max-interval <integer>
set ip6-min-interval <integer>
set ip6-other-flag <enable | disable>
set ip6-reachable-time <integer>
set ip6-retrans-time <integer>
  config ip6-extra-addr
    edit <ipv6_addr>
  end
  config ip6-prefix-list
    set autonomous-flag <enable | disable>
    set onlink-flag <enable | disable>
    set preferred-life-time <integer>
    set valid-life-time <integer>
  end
end
end
end

```

Configuring IPv6 routing

IPv6 routing is supported in both static and dynamic routing. The main difference from a configuration point of view is in the addresses.

Static routing

Static routing for IPv6 is essentially the same as with IPv4. From a configuration point of view, the only difference is the type of addresses used. When both IPv4 and IPv6 static routes are configured, they are displayed under two separate headings on the static routing page - *Route* and *IPv6 Route*. Use the arrows next to each heading to expand or minimize that list of routes.

To configure IPv6 static routes - web-based manager

- 1 Go to *Router > Static > Static Route*.
- 2 Select arrow to expand the *Create New* menu.
- 3 Select *IPv6 Route*.
- 4 Enter *Destination IP/Mask*, *Device*, *Gateway*, *Distance*, and *Priority* as with normal static routing using IPv6 addresses.
- 5 Select *OK*.

To configure IPv6 static routes - CLI

Use the following command to add an IPv6 static route:

```

config router static6
  edit 1
    set dst <ipv6_addr>
    set gateway <ipv6_addr>
    set device <interface>
  end
end

```

```

    set priority <integer>
end

```

Dynamic routing

As with static routing, the dynamic routing protocols all have IPv6 versions. Both IPv4 and IPv6 dynamic routing can be running at the same time due to the dual stack architecture of the FortiGate unit. IPv6 dynamic routing must be configured using CLI commands.

Table 7: Dynamic routing protocols, IPv6 versions, CLI command, and RFCs

Dynamic Routing	IPv6	CLI command	IPv6 RFC
RIP	RIP next generation (RIPng)	config router ripng	RFC 2080
BGP	BGP4+	config router bgp All parts of bgp that include IP addresses have IPv4 and IPv6 versions.	RFC 2545 and RFC 2858
OSPF	OSPFv3	config router ospf6	RFC 2740

Configuring IPv6 security policies

Configuring IPv6 security policies is similar to configuring IPv4 security policies. On the web-based manager go to *Policy > IPv6 Policy*. From the CLI use the command `config firewall policy6`. You must also add IPv6 firewall addresses (*Firewall Objects > Address* or `config firewall address6`) and address groups (*Firewall Objects > Address > Group* or `config firewall addgrp6`).

Under the security policies for IPv6, you can also define SSL-VPN actions and authentication policies.

IPv6 Policy configuration settings

The following are IPv6 security policy configuration settings in *Policy > Policy > IPv6 Policy*.

New Policy page	
Source Interface/Zone	Select the name of the FortiGate network interface, virtual domain (VDM) link, or zone on which IP packets are received. Interfaces and zones are configured on the System Network page. You can also create a web proxy firewall proxy by selecting <i>web-proxy</i> in Source Interface/Zone.
	If you select <i>any</i> as the source interface, the security policy matches all interfaces as source. When you select <i>any</i> as the source interface, that security policy list is displayed only in global view.
	If <i>Action</i> is set to <i>IPSEC</i> , the interface is associated with the local private network.
	If <i>Action</i> is set to <i>SSL-VPN</i> , the interface is associated with connections from remote SSL VPN clients.
Source Address	Select the name of a firewall address to associate with the <i>Source Interface/Zone</i> .

		<p>You can also create firewall addresses by selecting <i>Create New</i> from this list.</p> <p>If you want to associate multiple firewall addresses or address groups with <i>Source Interface/Zone</i>, from <i>Source Address</i>, select <i>Multiple</i>. In the dialog box, move the firewall addresses or address groups from the <i>Available Addresses</i> section to the <i>Members</i> section, then select <i>OK</i>.</p>
Destination Interface/Zone		Select the name of the FortiGate network interface, virtual domain (VDM) link, or zone to which IP packets are forwarded. Interfaces and zones are configured on the System Network page.
		If you select <i>any</i> as the source interface, the security policy matches all interfaces as source. When you select <i>any</i> as the source interface, that security policy list is displayed only in <i>Global View</i> .
Destination Address		Select the name of a firewall address to associate with <i>Destination Interface/Zone</i> . Only packets whose header contains an IP address matching the selected firewall address will be subject to this security policy.
		You can also create firewall addresses by selecting <i>Create New</i> from this list.
		If you want to associate multiple firewall addresses or address groups with the <i>Destination Interface/Zone</i> , from <i>Destination Address</i> , select <i>Multiple</i> . In the dialog box, move the firewall addresses or address groups from the <i>Available Addresses</i> section to the <i>Members</i> section, then select <i>OK</i> .
		If you select a virtual IP, the unit applies NAT or PAT. The applied translation varies by the settings specified in the virtual IP, and whether you select NAT (below).
Schedule		<p>Select a one-time or recurring schedule or a schedule group that controls when the security policy is in effect.</p> <p>You can also create schedules by selecting <i>Create New</i> from this list.</p>
Service		<p>Select a firewall service or create a new custom service.</p> <p>If you are creating a web proxy security policy, <i>Web Proxy Service</i> appears and you can choose either a web proxy service or web proxy group.</p>
Action		Select how you want the firewall to respond when a packet matches the conditions of the security policy.
Log Allowed Traffic		<p>Select to record security policy traffic activity whenever the security policy processes a connection. These log messages are located in the traffic log.</p> <p>You must also enable traffic log for a logging location and set the logging severity level to <i>Notification</i> or lower using the Log&Report menu.</p> <p>This option is not available for web-proxy security policies.</p>
	Log Violation traffic	<p>Select to record security policy traffic activity whenever the security policy processes a violation. These log messages are located in the traffic log.</p> <p>Appears only when <i>Action</i> is <i>DENY</i>.</p>

Enable web cache	<p>Select to enable web caching for HTTP traffic accepted by the security policy. This option is available only on FortiGate units that support WAN Optimization and web caching. Enabling web caching in a security policy is similar to enabling web caching in a WAN Optimization rule. However, enabling web caching in a security policy means you can also apply UTM options to web cached traffic in a single VDOM.</p> <p>You can use this option to apply web caching for explicit web proxy traffic if the Source Interface/Zone is set to the web-proxy interface.</p> <p>Web caching supports caching of HTTP 1.0 and HTTP 1.1 web sites on the FortiGate unit hard disk. Some HTTP content accepted by the security policy may not be cached. See RFC 2616 for information about web caching for HTTP 1.1.</p>
Enable NAT	<p>Available only if <i>Action</i> is set to <i>ACCEPT</i> or <i>SSL-VPN</i>. Enable or disable Network Address Translation (NAT) of the source address and port of packets accepted by the security policy. When <i>NAT</i> is enabled, you can also configure <i>Dynamic IP Pool</i> and <i>Fixed Port</i>.</p> <p>If you select a virtual IP as the <i>Destination Address</i>, but do not select the <i>NAT</i> option, the unit performs destination NAT (DNAT) rather than full NAT. Source NAT (SNAT) is not performed.</p>
Use Destination Interface Address	<p>Select to use the destination interface address. If <i>Central NAT Table</i> is enabled, you can choose between this option and using the central NAT table.</p>
Use Central NAT Table	<p>Select to enabling logging using the Central NAT table that you configured in the Central NAT Table menu.</p>
Use Dynamic IP Pool	<p>Available only when <i>Enable NAT</i> is selected.</p> <p>Select the check box, then select an IP pool to translate the source address to an IP address randomly selected from addresses in the IP Pool.</p> <p><i>IP Pool</i> cannot be selected if the destination interface, VLAN subinterface, or one of the interfaces or VLAN subinterfaces in the destination zone is configured using DHCP or PPPoE.</p>
Enable Identity Based Policy	<p>Select to configure security policies that require authentication.</p>
Resolve User Names Using FSSO Agent	<p>Select to resolve user names when using the Fortinet Single Sign-On Agent feature.</p>

Enable Dynamic Profile	<p>Select to configure a dynamic profile security policy. Dynamic profile is a method for users to use a RADIUS server for a single sign-on access to network resources.</p> <p>The <i>Enable Dynamic Profile</i> option does not display by default on the web-based manager; you must first enable it in <i>System > Admin > Settings</i>. If you have VDOMs enabled, you can configure one RADIUS server and security policy for dynamic profile per VDOM. With multiple VDOMs, you can have each one with their own profile group on their own RADIUS server with their own custom level of access.</p> <p>After selecting the check box beside <i>Enable Dynamic Profile</i>, the following options appear below:</p> <ul style="list-style-type: none"> • <i>Profile Group</i> – select a dynamic profile group from the drop-down list • <i>Dynamic Profile Users Only</i> – select to only accept sessions with source addresses that are in the user context list
UTM	Select an UTM option to apply to the security policy. You must enable UTM before you can select the available UTM options. When selecting an option, select a profile from the list, or select <i>Create New</i> from the list to build a profile.
Web Proxy Forwarding Server	Select a web proxy forwarding server from the drop-down list. This appears only when configuring a web proxy security policy.
GTP Profile (FortiOS Carrier only)	Select a GTP profile from the drop-down list. Select <i>Create New</i> to create a new GTP profile. Select <i>View</i> to view the GTP profile.
Traffic Shaping	Select a traffic shaper for the security policy. You can also create a new shared traffic shaper. Shared traffic shapers control the bandwidth available to and set the priority of the traffic as its processed by, the security policy.
Reverse Direction Traffic Shaping	Select to enable reverse traffic shaping and select a shared traffic shaper. For example, if the traffic direction that a security policy controls is from port1 to port2, select this option will also apply the security policy shaping configuration to traffic from port2 to port1.
Dynamic Profile Users Only	Select to configure the security policy to only accept sessions with source addresses that are in the dynamic profile user context list. Sessions with source addresses that are not in the user context list do not match the security policy. For sessions that do not match the security policy, the unit continues searching down the security policy list for a match.
Enable Endpoint Security	<p>Select to enable the Endpoint NAC feature and select the Endpoint NAC profile to apply.</p> <ul style="list-style-type: none"> • You cannot enable Endpoint in security policies if <i>Redirect HTTP Challenge to a Secure Channel (HTTPS)</i> is enabled in <i>User > Options > Authentication</i>. • If the security policy involves a load balancing virtual IP, the Endpoint check is not performed.
Enable Disclaimer	Select to include a disclaimer page. Select <i>Edit</i> to modify the disclaimer replacement message.

Tags	Applies tags to the security policy. Tags can be viewed on the Policy page in the <i>Tags</i> column.
Applied Tags	Displays the tags that you have added to the security policy.
Add Tags	Enter the tag in the field and select the plus (+) sign to add the tag to the security policy. This also adds the tag to the <i>Applied Tags</i> list.
Comments	Add information about the security policy. The maximum length is 63 characters.

Configuring IPv6 DNS

Configuring DNS servers with IPv6 addresses is located in the same location as IPv4, by going to *System > Network > DNS*. There is a separate area for adding IPv6 addresses for DNS. From the CLI, use the command `config system dns`, where additional commands `ip6-primary` and `ip6-secondary` are available.

Configuring IPv6 DHCP

Configuring DHCP servers with IPv6 is performed using the CLI only. While similar to IPv4, there are a few exceptions:

- There is no gateway to define. A host learns the gateway using router advertisement messages
- There is no WINS servers defined for dhcpv6, as it is obsolete.

To configure DHCP use the following command set:

```
config system dhcp6 server
  edit 1
    set domain example.com
    set interface port3
    config ip-range
      edit 1
        set end-ip 2800:68:15:3::10
        set start-ip 2800:68:15:3::1
      end
    set option1 50 'AABB'
    set subnet 2800:68:15:3::/64
    set dns-server1 2800:68:15:3::2
    set dns-server2 2800:68:15:3::29
    set dns-server3 2800:68:15:3::28
    set enable enable
  end
end
```

For more information on the commands, see the [CLI Reference](#).

Configuring IPv6 over IPv4 tunneling

IPv6 over IPv4 tunneling can only be configured in the CLI using the `config system sit-tunnel` command. When you configure an IPv6 over IPv4 tunnel, you are creating a virtual interface that can be used in configurations just like any other virtual interface such as VLANs.

The name of the command `sit-tunnel` comes from Simple Internet Transition (SIT) tunneling. For the period while IPv6 hosts and routers co-exist with IPv4, a number of transition mechanisms are needed to enable IPv6-only hosts to reach IPv4 services and to allow isolated IPv6 hosts and networks to reach the IPv6 Internet over the IPv4 infrastructure.

These techniques, collectively called Simple Internet Transition, include:

- dual-stack IP implementations for interoperating hosts and routers
- embedding IPv4 addresses in IPv6 addresses
- IPv6-over-IPv4 tunneling mechanisms
- IPv4/IPv6 header translation

The syntax for the IPv6 over IPv4 tunneling CLI command is:

```
config system sit-tunnel
  edit <name_string>
    set destination <ipv4_addr>
    set interface <interface_string>
    set ip6 <ipv6_addr>
    set source <ipv4_addr>
  next
end
```

<name_string>	This will be the name of the tunnel, and appear in the network interface list. It should be descriptive such as <code>my_ip6_tunnel</code> . The maximum length allowed is 15 characters.
destination <ipv4_addr>	This is the tunnel broker's IPv4 server address. It is one of the two ends of the tunnel.
interface <interface_string>	This interface is the interface the tunnel piggy backs on. Generally this should be the external interface of the FortiGate unit. This setting is optional if you don't have a fixed IP address from your ISP.
ip6 <ipv6_addr>	The IPv6 address of the tunnel.
source <ipv4_addr>	This is the FortiGate unit end of the tunnel. It is just like any other FortiGate unit interface address. If this address is DHCP-based, it will change. In that case you should ensure the netmask covers the possible range of addresses. It is possible to use 0.0.0.0 to cover all possible addresses if you have a DDNS or PPPoE connection where the address changes.

For more configuration of tunnels see [“Configuring FortiOS to connect to an IPv6 tunnel provider” on page 74](#).

Configuring IPv6 IPsec VPNs

The FortiGate unit supports route-based IPv6 IPsec, but not policy-based.

Where both the gateways and the protected networks use IPv6 addresses, sometimes called IPv6 over IPv6, you can create either an auto-keyed or manually-keyed VPN. You can combine IPv6 and IPv4 addressing in an auto-keyed VPN in the following ways:

IPv4 over IPv6	<p>The VPN gateways have IPv6 addresses.</p> <p>The protected networks have IPv4 addresses. The phase 2 configurations at either end use IPv4 selectors.</p>
IPv6 over IPv4	<p>The VPN gateways have IPv4 addresses.</p> <p>The protected networks use IPv6 addresses. The phase 2 configurations at either end use IPv6 selectors.</p>

Compared with IPv4 IPsec VPN functionality, there are some limitations:

- Except for IPv6 over IPv4, remote gateways with Dynamic DNS are not supported. This is because FortiOS 3.0 does not support IPv6 DNS.
- You cannot use RSA certificates in which the common name (cn) is a domain name that resolves to an IPv6 address. This is because FortiOS 3.0 does not support IPv6 DNS.
- DHCP over IPsec is not supported, because FortiOS 3.0 does not support IPv6 DHCP.
- Selectors cannot be firewall address names. Only IP address, address range and subnet are supported.
- Redundant IPv6 tunnels are not supported.

Certificates

On a VPN with IPv6 phase 1 configuration, you can authenticate using VPN certificates in which the common name (cn) is an IPv6 address. The `cn-type` keyword of the `user peer` command has an option, `ipv6`, to support this.

Configuring IPv6 IPsec VPNs

Configuration of an IPv6 IPsec VPN follows the same sequence as for an IPv4 route-based VPN: phase 1 settings, phase 2 settings, security policies and routing.

To access IPv6 functionality through the web-based manager, go to *System Admin > Settings* and enable *IPv6 Support on GUI*.

Phase 1 configuration

In the web-based manager, you define the Phase 1 as IPv6 in the *Advanced* settings. Enable the *IPv6 Version* check box. You can then enter an IPv6 address for the remote gateway.

In the CLI, you define an IPsec phase 1 configuration as IPv6 by setting `ip-version` to 6. Its default value is 4. Then, the `local-gw` and `remote-gw` keywords are hidden and the corresponding `local-gw6` and `remote-gw6` keywords are available. The values for `local-gw6` and `remote-gw6` must be IPv6 addresses.

For example:

```
config vpn ipsec phase1-interface
edit tunnel6
set ip-version 6
set remote-gw6 0:123:4567::1234
set interface port3
set proposal 3des-md5
end
```

Phase 2 configuration

To create an IPv6 IPsec phase 2 configuration in the web-based manager, you need to define IPv6 selectors in the Advanced settings. Change the default 0.0.0.0/0 address for Source address and Destination address to the IPv6 value ::/0. If needed, enter specific IPv6 addresses, address ranges or subnet addresses in these fields.

In the CLI, set `src-addr-type` and `dst-addr-type` to `ip6`, `range6` or `subnet6` to specify IPv6 selectors. By default, zero selectors are entered, ::/0 for the `subnet6` address type, for example. The simplest IPv6 phase 2 configuration looks like the following:

```
config vpn ipsec phase2-interface
edit tunnel6_p2
set phase1name tunnel6
set proposal 3des-md5
set src-addr-type subnet6
set dst-addr-type subnet6
end
```

Security policies

To complete the VPN configuration, you need a security policy in each direction to permit traffic between the protected network's port and the IPsec interface. You need IPv6 policies unless the VPN is IPv4 over IPv6.

Routing

Appropriate routing is needed for both the IPsec packets and the encapsulated traffic within them. You need a route, which could be the default route, to the remote VPN gateway via the appropriate interface. You also need a route to the remote protected network via the IPsec interface.

To create a static route in the web-based manager, go to *Router > Static > Static Route*. Select the drop-down arrow for *Create New* and select *IPv6 Route*. Enter the information and select *OK*. In the CLI, use the `router static6` command. For example, where the remote network is `fec0:0000:0000:0004::/64` and the IPsec interface is `toB`:

```
config router static6
edit 1
set device port2
set dst 0::/0
next
edit 2
set device toB
set dst fec0:0000:0000:0004::/64
next
end
```

If the VPN is IPv4 over IPv6, the route to the remote protected network is an IPv4 route. If the VPN is IPv6 over IPv4, the route to the remote VPN gateway is an IPv4 route.

IPv6 troubleshooting

There are a number of troubleshooting methods that can be used with IPv6 issues.

ping6

The main method of troubleshooting IPv6 traffic is using the IPv6 version of ping.

You can use the IPv6 ping command to:

- send an ICMP echo request packet to the IPv6 address that you specify.
- specify a source interface other than the one from which the probe originates by using the source interface keywords.
- specify a source IP address other than the one from which the probe originates by using the source address keywords

You can specify the following options:

packetCount	Number of packets to send to the destination IPv6 address. If you specify a zero, echo requests packets are sent indefinitely.
data-pattern	Sets the type of bits contained in the packet to all ones, all zeros, a random mixture of ones and zeros, or a specific hexadecimal data pattern that can range from 0x0 to 0xFFFFFFFF. The default is all zeros.
extended header attributes	Set the interface type and specifier of a destination address on the system that is configured for external loopback; the command succeeds only if the specified interface is configured for external loopback.
sweep interval	Specifies the change in the size of subsequent ping packets while sweeping across a range of sizes. For example, you can configure the sweep interval to sweep across the range of packets from 100 bytes to 1000 bytes in increments specified by the sweep interval. By default, the system increments packets by one byte; for example, it sends 100, 101, 102, 103, ... 1000. If the sweep interval is 5, the system sends 100, 105, 110, 115, ... 1000.
sweep sizes	Enables you to vary the sizes of the echo packets being sent. Used to determine the minimum sizes of the MTUs configured on the nodes along the path to the destination address. This reduces packet fragmentation, which contributes to performance problems. The default is to not sweep (all packets are the same size).
timeout	Sets the number of seconds to wait for an ICMP echo reply packet before the connection attempt times out.
hop limit	Sets the time-to-live hop count in the range 1-255; the default is 255.

The following characters may appear in the display after the ping command is issued:

- ! - reply received
- . - timed out while waiting for a reply
- ? - unknown packet type
- A - admin unreachable
- b - packet too big
- H - host unreachable
- N - network unreachable
- P - port unreachable
- p - parameter problem
- S - source beyond scope

t - hop limit expired (TTL expired)

IPv6 ping description

Ping uses the ICMP protocol's mandatory ECHO_REQUEST datagram to elicit an ICMP ECHO_RESPONSE from a host or gateway. ECHO_REQUEST datagrams ("pings") have an IP and ICMP header, followed by a strict timeval and then an arbitrary number of "pad" bytes used to fill out the packet.

See also

IPv6 ping options

-a	Audible ping.
-A	Adaptive ping. Interpacket interval adapts to round-trip time, so effectively no more than one (or more, if preload is set) unanswered probe is present in the network. Minimal interval is 200msec for any user other than administrator. On networks with low rtt this mode is essentially equivalent to flood mode.
-b	Allow pingging of a broadcast address.
-B	Do not allow ping to change source address of probes. The address is bound to one selected when the ping starts.
-c count	Stop after sending count ECHO_REQUEST packets. With deadline option, ping waits for count ECHO_REPLY packets, until the timeout expires.
-d	Set the SO_DEBUG option on the socket being used. This socket option is not used by a Linux kernel.
-F flow label	Allocate and set 20 bit flow label on echo request packets (only ping6). If value is zero, kernel allocates random flow label.
-f	Flood ping. For every ECHO_REQUEST sent a period "." is displayed, while for ever ECHO_REPLY received a backspace is displayed. This provides a rapid display of how many packets are being dropped. If interval is not specified, it is set to zero and packets are output as fast as they come back or one hundred times per second, whichever is faster. Only the administrator may use this option with zero interval.
-i interval	Wait a specified interval of seconds between sending each packet. The default is 1 second between each packet, or no wait in flood mode. Only an administrator can set the interval to a value of less than 0.2 seconds.
-I interface address	Set source address to specified interface address. Argument may be numeric IP address or name of device. This option is required when you ping an IPv6 link-local address.
-l preload	If preload is specified, ping sends this number of packets that are not waiting for a reply. Only the administrator may select a preload of more than 3.
-L	Suppress loopback of multicast packets. This flag only applies if the ping destination is a multicast address.
-n	Numeric output only. No attempt will be made to look up symbolic names for host addresses.

-p pattern	You may specify up to 16 "pad" bytes to fill out the packet you send. This is useful for diagnosing data-dependent problems in a network. For example, -p ff will cause the sent packet to be filled with all ones.
-Q tos	Set Quality of Service -related bits in ICMP datagrams. tos can be either decimal or hex number. Traditionally (RFC1349), these have been interpreted as: 0 for reserved (currently being redefined as congestion control), 1-4 for Type of Service and 5-7 for Precedence. Possible settings for Type of Service are: minimal cost: 0x02, reliability: 0x04, throughput: 0x08, low delay: 0x10. Multiple TOS bits should not be set simultaneously. Possible settings for special Precedence range from priority (0x20) to net control (0xe0). You must be root (CAP_NET_ADMIN capability) to use Critical or higher precedence value. You cannot set bit 0x01 (reserved) unless ECN has been enabled in the kernel. In RFC 2474, these fields has been redefined as 8-bit Differentiated Services (DS), consisting of: bits 0-1 of separate data (ECN will be used, here), and bits 2-7 of Differentiated Services Codepoint (DSCP).
-q	Quiet output. Nothing is displayed except the summary lines at startup time and when finished
-R	Record route. (IPv4 only) Includes the RECORD_ROUTE option in the ECHO_REQUEST packet and displays the route buffer on returned packets. Note that the IP header is only large enough for nine such routes. Many hosts ignore or discard this option.
-r	Bypass the normal routing tables and send directly to a host on an attached interface. If the host is not on a directly-attached network, an error is returned. This option can be used to ping a local host through an interface that has no route through it provided the option -I is also used.
-s packetsize	Specifies the number of data bytes to be sent. The default is 56, which translates into 64 ICMP data bytes when combined with the 8 bytes of ICMP header data.
-S sndbuf	Set socket sndbuf (send buffer). If not specified, it is selected to buffer not more than one packet.
-t ttl	Set the IP Time to Live.
-T timestamp option	Set special IP timestamp options. May be either tsonly (only timestamps), tsandaddr (timestamps and addresses) or tsprespec host1 [host2 [host3 [host4]]] (timestamp prespecified hops).
-M hint	Select Path MTU Discovery strategy. hint may be either do (prohibit fragmentation, even local one), want (do PMTU discovery, fragment locally when packet size is large), or don't (do not set DF flag).
-U	Print full user-to-user latency (the old behavior). Normally ping prints network round trip time, which can be different f.e. due to DNS failures.
-v	Verbose output.
-V	Show version and exit.

-w deadline	Specify a timeout, in seconds, before ping exits regardless of how many packets have been sent or received. In this case ping does not stop after count packet are sent, it waits either for deadline expire or until count probes are answered or for some error notification from network.
-W timeout	Time to wait for a response, in seconds. The option affects only timeout in absence of any responses, otherwise ping waits for two RTTs.

Examples

Ping a global V6 address with a 1400 byte packet from FortiGate CLI:

```
execute ping6 -s 1400 2001:480:332::10
```

Ping a multicast group using a ping6 command on FortiGate CLI (-I and port name must be specified for CLI ping6 command to ping v6 multicast group):

```
execute ping6 -I port1 ff02::1
```

Ping a localnet v6 address from FortiGate CLI:

```
execute ping6 FE80:0:0:0:213:e8ff:fe9e:ccf7
```

This address would normally be written as FE80::213:e8ff:fe9e:ccf7.

diagnose sniffer packet

The FortiOS built in packet sniffer also works with IPv6. The following are some examples using an IPv6-over-IPv4 tunnel called test6.

```
diagnose sniffer packet test6 'none' 4
interfaces=[test6]
filters=[]
pcap_lookupnet: test6: no IPv4 address assigned
34.258651 test6 -- 2001:4dd0:ff00:15d::2 -> 2001:4dd0:ff00:15d::1:
icmp6: echo request seq 1
34.324658 test6 -- 2001:4dd0:ff00:15d::1 -> 2001:4dd0:ff00:15d::2:
icmp6: echo reply seq 1
35.268581 test6 -- 2001:4dd0:ff00:15d::2 -> 2001:4dd0:ff00:15d::1:
icmp6: echo request seq 2
35.334230 test6 -- 2001:4dd0:ff00:15d::1 -> 2001:4dd0:ff00:15d::2:
icmp6: echo reply seq
```

```
diagnose sniffer packet any 'ip6 and tcp port 80' 4 10
interfaces=[any]
filters=[ip6 and tcp port 80]
1 LAN in 2001:4dd0:ff42:72:21b:63ff:fe08:e071.53037 ->
2a00:1450:8007::63.80: syn 2298823882
2 test6 out 2001:4dd0:ff42:72:21b:63ff:fe08:e071.53037 ->
2a00:1450:8007::63.80: syn 2298823882
3 test6 in 2a00:1450:8007::63.80 ->
2001:4dd0:ff42:72:21b:63ff:fe08:e071.53037: syn 4218782319
ack
4 LAN out 2a00:1450:8007::63.80 ->
2001:4dd0:ff42:72:21b:63ff:fe08:e071.53037: syn 4218782319
ack
5 LAN in 2001:4dd0:ff42:72:21b:63ff:fe08:e071.53037 ->
2a00:1450:8007::63.80: ack 4218782320
```

```
6 test6 out 2001:4dd0:ff42:72:21b:63ff:fe08:e071.53037 ->
2a00:1450:8007::63.80: ack 4218782320
```

diagnose debug flow

The `diagnose debug flow` CLI command is the same for IPv6 or IPv4. The output format is the same, however the command is only slightly different in that it uses `filter6` and an IPv6 address.

To enable diag debug flow for IPv6 - CLI

```
# diagnose debug enable
# diagnose debug flow show console enable
# diagnose debug flow show func enable
# diagnose debug flow filter6 addr 2001:4dd0:ff42:12::24
# diagnose debug flow trace start6
```

IPv6 specific diag commands

To list all the sit-tunnels that are configured:

```
diagnose ipv6 sit-tunnel list
total tunnel = 1:
devname=test6 devindex=4 ifindex=22 saddr=0.0.0.0
daddr=88.25.29.134 proto=41 vfid=0000 ref=2
```

To list all the IPv6 routes:

```
diagnose ipv6 route list
vf=0 type=02 protocol=unspec flag=00200001 oif=8(root)
dst::1/128 gwy:: prio=0
vf=0 type=02 protocol=unspec flag=00200001 oif=8(root)
dst:2001:4dd0:ff00:75d::2/128 gwy:: prio=0
vf=0 type=01 protocol=kernel flag=00240021 oif=22(sixxs)
dst:2001:4dd0:ff00:75d::/64 gwy:: prio=100
vf=0 type=02 protocol=unspec flag=00200001 oif=8(root)
dst:2001:4dd0:ff42:68::1/128 gwy:: prio=0
vf=0 type=01 protocol=kernel flag=01040001 oif=19(LAN)
dst:2001:4dd0:ff42:68:225:ff:feee:5314/128
gwy:2001:4dd0:ff42:68:225:ff:feee:5314 prio=0
.....
```


Some other IPv6 diagnose commands include:

<code>diagnose ipv6 neighbor-cache</code>	Add, delete, flush, or list the IPv6 ARP table or table entry.
<code>diagnose sys session6</code>	Clear, filter, full-stat, list, stat IPv6 sessions.
<code>tree diagnose ipv6</code>	View all the diagnose IPv6 commands.

Additional IPv6 resources

There are many RFCs available regarding IPv6. The following table lists the major IPv6 articles and their Internet Engineering Task Force (IETF) web locations.

RFC	Subject	Location
RFC 1933, <i>Transition Mechanisms for IPv6 Hosts and Routers</i>	Describes IPv4 compatibility mechanisms that can be implemented by IPv6 hosts and routers	http://www.ietf.org/rfc/rfc1933
RFC 2185, <i>Routing Aspects of IPv6 Transition</i>	Provides an overview of the routing aspects of the IPv6 transition	http://www.ietf.org/rfc/rfc2185
RFC 2373, <i>IP Version 6 Addressing Architecture</i>	Defines the addressing architecture of the IP Version 6 protocol [IPv6]	http://www.ietf.org/rfc/rfc2373
RFC 2402, <i>IP Authentication Header</i>	Describes functionality and implementation of IP Authentication Headers (AH)	http://www.ietf.org/rfc/rfc2402
RFC 2460, <i>Internet Protocol, Version 6 (IPv6) Specification</i>	Describes functionality, configuration of IP version 6 (IPv6) and differences from IPv4.	http://www.ietf.org/rfc/rfc2460
RFC 2461, <i>Neighbor Discovery for IP Version 6 (IPv6)</i>	Describes the features and functions of IPv6 Neighbor Discovery protocol	http://www.ietf.org/rfc/rfc2461
RFC 2462, <i>IPv6 Stateless Address Autoconfiguration</i>	Specifies the steps a host takes in deciding how to autoconfigure its interfaces in IPv6	http://www.ietf.org/rfc/rfc2462
RFC 2893, <i>Transition Mechanisms for IPv6 Hosts and Routers</i>	Specifies IPv4 compatibility mechanisms that can be implemented by IPv6 hosts and routers	http://www.ietf.org/rfc/rfc2893
RFC 3306, <i>Unicast-Prefix-Based IPv6 Multicast Addresses</i>	Describes the format and types of Ipv6 multicast addresses	http://www.ietf.org/rfc/rfc3306
RFC 3484, <i>Default Address Selection for Internet protocol version 6 (IPv6)</i>	Describes the algorithms used in IPv6 default address selection	http://www.ietf.org/rfc/rfc3484
RFC 3513, <i>Internet Protocol version 6 (IPv6) Addressing Architecture</i>	Contains details about the types of IPv6 addresses and includes examples	http://www.ietf.org/rfc/rfc3513
RFC 3587, <i>IPv6 Global Unicast Address Format</i>	Defines the standard format for IPv6 unicast addresses	http://www.ietf.org/rfc/rfc3587



Advanced FortiGate firewall concepts

The FortiGate firewall has advanced firewall component options, which allows for greater flexibility when these advanced options are needed to help with your growing network. These advanced firewall components include traffic shaping, QoS and identity-based policies.

The following topics are included in this section:

- [Central NAT table](#)
- [Stateful inspection of SCTP traffic](#)
- [Port pairing](#)
- [Blocking port 25 to email server traffic](#)
- [Blocking HTTP access by IP](#)
- [ICMP packet processing](#)
- [Adding NAT security policies in Transparent mode](#)
- [Adding a static NAT virtual IP for a single IP address and port](#)
- [Double NAT: combining IP pool with virtual IP](#)
- [Using VIP range for Source NAT \(SNAT\) and static 1-to-1 mapping](#)
- [Traffic shaping and per-IP traffic shaping](#)
- [Endpoint Security](#)
- [Logging traffic](#)
- [Quality of Service \(QoS\)](#)
- [Identity-based security policies](#)

Central NAT table

The central NAT table enables you to define, and control with more granularity, the address translation performed by the FortiGate unit. With the NAT table, you can define the rules which dictate the source address or address group and which IP pool the destination address uses.

While similar in functionality to IP pools, where a single address is translated to an alternate address from a range of IP addresses, with IP pools there is no control over the translated port. When using the IP pool for source NAT, you can define a fixed port to guarantee the source port number is unchanged. If no fix port is defined, the port translation is randomly chosen by the FortiGate unit. With the central NAT table, you have full control over both the IP address and port translation.

The NAT table also functions in the same way as the security policy table. That is, the FortiGate unit reads the NAT rules in a top-down methodology, until it hits a matching rule for the incoming address. This enables you to create multiple NAT policies that dictate which IP pool is used based on the source address. The NAT policies can be rearranged within the policy list as well, the same way as security policies. NAT policies are applied to network traffic after a security policy.

To view the Central NAT configuration page, and use them in a security policy, you need to first enable it.

To enable Central NAT - web-based manager

- 1 Go to *System > Admin > Settings*.
- 2 In the Display Options on GUI section, select the check box beside *Central NAT table*.
- 3 Select *Apply*.

To enable Central NAT - CLI

```
config system global
    set gui-central-nat-table
end
```

NAT policies are created in the web-based manager by going to *Policy > Policy > Central NAT Table*. The NAT policies are enabled when you configure the security policy by selecting the *Use Central NAT Table* option.

NAT policies are created in the CLI by using the commands under `config firewall central-nat`. To enable the policies use the commands

```
config security policy
    edit <policy_number>
        set central-nat enable
    end
```

Central NAT Table configuration settings

To configure the Central NAT table, go to *Policy > Policy > Central NAT Table* and select *Create New*.

New NAT page

Source Address	Select the source IP address from the drop-down list. You can optionally create a group of source IP addresses when you select <i>Multiple</i> in the drop-down list. You can also create a new source IP address when you select <i>Create New</i> in the drop-down list.
Translated Address	Select the dynamic IP pool from the drop-down list.
Original Source Port	Enter the source port that the address is originating from.
Translated Port	Enter the translated port number. The number in the <i>From</i> field must be greater than the lower port number that is entered in the <i>To</i> field.

Stateful inspection of SCTP traffic

The Stream Control Transmission Protocol (SCTP) is a transport layer protocol similar to TCP and UDP. SCTP is designed to provide reliable, in-sequence transport of messages with congestion control. SCTP is defined in [RFC 4960](http://tools.ietf.org/html/rfc4960).

Some common applications of SCTP include supporting transmission of the following protocols over IP networks:

- SCTP is important in 3G and 4G/LTE networks (for example, HomeNodeB = FemtoCells)
- SS7 over IP (for example, for 3G mobile networks)
- SCTP is also defined and used for SIP over SCTP and H.248 over SCTP
- Transport of Public Switched Telephone Network (PSTN) signaling messages over IP networks.

SCTP is a reliable transport protocol that runs on top of a connectionless packet network (IP). SCTP provides the following services:

- Acknowledged error-free non-duplicated transfer of user data
- Data fragmentation to conform to discovered path MTU size
- Sequenced delivery of user messages within multiple streams, with an option for order-of-arrival delivery of individual user messages
- Optional bundling of multiple user messages into a single SCTP packet
- network-level fault tolerance through supporting of multi-homing at either or both ends of an association
- Congestion avoidance behavior and resistance to flooding and masquerade attacks

SCTP is effective as the transport protocol for applications that require monitoring and session-loss detection. For such applications, the SCTP path and session failure-detection mechanisms actively monitor the connectivity of the session. SCTP differs from TCP in having multi-homing capabilities at either or both ends and several streams within a connection, typically referred to as an association. A TCP stream represents a sequence of bytes; an SCTP stream represents a sequence of messages.

Configuring FortiGate SCTP filtering

The FortiGate firewall can apply security policies to SCTP sessions in the same way as TCP and UDP sessions. You can create security policies that accept or deny SCTP traffic by setting the service to ANY. FortiOS does not include pre-defined SCTP services. To configure security policies for traffic with specific SCTP source or destination ports you must create custom firewall services for SCTP.

FortiGate units route SCTP traffic in the same way as TCP and UDP traffic. You can configure policy routes specifically for routing SCTP traffic by setting the protocol number to 132. SCTP policy routes can route SCTP traffic according to the destination port of the traffic if you add a port range to the policy route.

You can configure a FortiGate unit to perform stateful inspection of different types of SCTP traffic by creating custom SCTP services and defining the port numbers or port ranges used by those services. FortiGate units support SCTP over IPv4. The FortiGate unit performs the following checks on SCTP packets:

- Source and Destination Port and Verification Tag.
- Chunk Type, Chunk Flags and Chunk Length
- Verify that association exists
- Sequence of Chunk Types (INIT, INIT ACK, etc)
- Timer checking
- Four way handshake checking
- Heartbeat mechanism

- Protection against INIT/ACK flood DoS attacks, and long-INIT flooding
- Protection against association hijacking

FortiOS also supports SCTP sessions over IPsec VPN tunnels, as well as full traffic and event logging for SCTP sessions.

Adding an SCTP custom service

This example creates a custom SCTP service that accepts SCTP traffic using destination port 2905. SCTP port number 2905 is used for SS7 Message Transfer Part 3 (MTP3) User Adaptation Layer (M3UA) over IP.

To add the SCTP custom service - web-based manager

- 1 Go to *Firewall Objects > Service > Custom* and select *Create New*.
- 2 Enter the following and select *OK*.

Name	M3UA_service
Protocol Type	TCP/UDP/SCTP
Protocol	SCTP
Source Port (Low)	1
Source Port (High)	65535
Destination Port (Low)	2905
Destination Port (High)	2905

To add the SCTP custom service - CLI

```
config firewall service custom
  edit M3UA_service
    set protocol TCP/UDP/SCTP
    set sctp-portrange 2905
  end
```

Adding an SCTP policy route

You can add policy routes that route SCTP traffic based on the SCTP source and destination port as well as other policy route criteria. The SCTP protocol number is 132.

The following example directs all SCTP traffic with SCTP destination port number 2905 to the next hop gateway at IP address 1.1.1.1.

To add the policy route - web-based manager

- 1 Go to *Router > Static > Policy Route*.
- 2 Select *Create New*.
- 3 Enter the following information and select *OK*.

Protocol	132
Incoming interface	internal
Source address / mask	0.0.0.0 0.0.0.0
Destination address / mask	0.0.0.0 0.0.0.0
Destination Ports	From 2905 to 2905

Force traffic to:

Outgoing interface	external
Gateway Address	1.1.1.1

To add the policy route - CLI

```
config router policy
  edit 1
    set input-device internal
    set src 0.0.0.0 0.0.0.0
    set dst 0.0.0.0 0.0.0.0
    set output-device external
    set gateway 1.1.1.1
    set protocol 132
    set start-port 2905
    set end-port 2905
  end
```

Changing the session time to live for SCTP traffic


Use the following command to change the session timeout for SCTP protocol M3UA on port 2905 to 3600 seconds.

```
config system session-ttl
  config port
    edit 1
      set protocol 132
      set start-port 2905
      set end-port 2905
      set timeout 3600
    end
  end
```

Port pairing

Port pairing is an option in Transparent mode to bind two ports together. In doing this, you can create security policies that regulate traffic only between two specific ports, VLANs or VDOMs. In its simplest form, this enables an administrator to create security policies that are only between these two ports. Traffic is captured between these ports. No other traffic can enter DNS services or leave a port pairing.

For example, a FortiGate unit has three ports, where port 1 and port 2 are paired together, because the two networks only need to communicate with each other. If packet arrives on port 1, the FortiGate unit needs to figure out whether the packet goes to port 2 or port 3. With port pairing configured, it is more simple. If packet arrives on port 1, then the FortiGate automatically directs the packet to port 2. The opposite is also true in the other direction. This can be ideal when to groups only need to transfer data between each other.

Figure 10: Port pairing**To configure port pairing - web-based manager**

- 1 Go to *System > Network > Interface*.
- 2 Select the arrow beside *Create New*, and select *Port Pair*.
- 3 Enter a *Name* for the port pair.
- 4 Select the physical or virtual ports from the *Available Members* list and select the right-facing arrow to add the ports to the *Selected Members* list.
There can be only two ports added.
- 5 Select *OK*.

To configure port pairing - CLI

```
config system port-pair
  edit <pair_name>
    set member <port_names>
  end
```

When configuring security policies with the port pairs, selecting the *Source Interface* automatically populates the *Destination Interface*, and vice versa. All other aspects of the security policy configuration remains the same.

Blocking port 25 to email server traffic

Port 25 is the default port for SMTP traffic. Certain types of malware can install themselves on an unsuspecting user's computer and send spam using its own email server. By blocking port 25, this prevents a host system, and potentially your network or company, from being deemed a spam source.

This does, however limit your corporation from using a web server. You have a few options for this:

- if the email server is on a dedicated port, such as a DMZ port, security policies can ensure no traffic goes out from this port except the email server.
- Block all traffic on port 25 except the specific address of the email server.

Dedicated traffic

This example shows the steps to ensure only traffic exits from the DMZ where the email server is connected. The internal port is connected to the internal network and the WAN1 port connects to the Internet.

First, create a security policy that will not allow any traffic through port 25 from the internal interface, which connects to the internal network. Place this policy at the top of the security policy list.

To block traffic on port 25 - web-based manager

- 1 Go to *Policy > Policy > Policy* and select *Create New*.
- 2 Set the following options and select *OK*.

Source Interface	Internal
Source Address	ALL
Destination Interface	WAN1
Destination Address	ALL
Schedule	ALWAYS
Service	SMTP
Action	DENY
Comments	Prevent Malware spam.

You may also want to enable *Log Violation Traffic* to see if there is any potential malware or other user sending email using the non-corporate email server.

To block traffic on port 25 - CLI

```
config security policy
edit <policy_number>
set srcintf Internal
set srcaddr all
set dstintf wan1
set dstaddr all
set schedule always
set service smtp
set action deny
set comment "Prevent Malware spam."
end
```

Next, create a security policy for the email server, IP address 10.10.11.29 that only allows SMTP traffic from the email server on port 25.

To allow traffic on port 25 for the email server - web-based manager

- 3 Go to *Policy > Policy > Policy* and select *Create New*.
- 4 Set the following options and select *OK*.

Source Interface	DMZ
Source Address	10.10.11.29
Destination Interface	WAN1
Destination Address	ALL

Schedule	ALWAYS
Service	SMTP
Action	ACCEPT

To allow traffic on port 25 for the email server- CLI

```
config security policy
  edit <policy_number>
    set srcintf dmz
    set srcaddr 10.10.11.29
    set dstintf wan1
    set dstaddr all
    set schedule always
    set service smtp
    set action allow
  end
```

Restricting traffic on port 25

This example shows how to limit traffic on port 25 on the wan port to only traffic from the email server. The web server's address is 10.10.10.29.

To allow traffic on port 25 for the email server - web-based manager

- 1 Go to *Policy > Policy > Policy* and select *Create New*.
- 2 Set the following options and select *OK*.

Source Interface	INTERNAL
Source Address	10.10.10.29
Destination Interface	WAN1
Destination Address	ALL
Schedule	ALWAYS
Service	SMTP
Action	ACCEPT

To allow traffic on port 25 for the email server- CLI

```
config security policy
  edit <policy_number>
    set srcintf internal
    set srcaddr 10.10.10.29
    set dstintf wan1
    set dstaddr all
    set schedule always
    set service smtp
    set action allow
  end
```

Next, add a deny security policy that blocks all SMTP traffic from the Internal port to the WAN1 port. Ensure this policy is directly after the policy created above.

To block SMTP traffic on port 25 for the rest of the company - web-based manager

- 3 Go to *Policy > Policy > Policy* and select *Create New*.

- 4 Set the following options and select *OK*.

Source Interface	INTERNAL
Source Address	ALL
Destination Interface	WAN1
Destination Address	ALL
Schedule	ALWAYS
Service	SMTP
Action	DENY

To block SMTP traffic on port 25 for the rest of the company - CLI

```
config security policy
  edit <policy_number>
    set srcintf internal
    set srcaddr all
    set dstintf wan1
    set dstaddr all
    set schedule always
    set service smtp
    set action deny
  end
```

Blocking HTTP access by IP

To block a web site using the IP, create a URL filter entry, using the additional information below. Note that this is only effective with HTTP or FortiGate units running Deep Inspection.

You need to create two URL filter entries. The first filter only allowing a text string containing two or more sets of text separated by a period. This is to match the various domain possibilities for web sites, for example:

- example.org
- www.example.com
- www.example.co.jp

The second filter blocks any IP address lookup.

To add the URL filter entries

- 1 Go to *UTM Profiles > Web Filter > URL Filter*.
- 2 Select *Create New* to add a filter group, give it a name and select *OK*.
- 3 Select *Create New* for a new filter.
- 4 Enter the *URL* of `^([a-z0-9-]+\.)\{1,\}[a-z]+\`
- 5 Set the *Type* to *Regex*.
- 6 Set the *Action* to *Allow*.
- 7 Select *OK*.
- 8 Select *Create New*.
- 9 Enter the *URL* of `[0-9]\{1,3\}\.[0-9]\{1,3\}\.[0-9]\{1,3\}\.[0-9]\{1,3\}`

10 Set the *Type* to *Regex*.

11 Set the *Action* to *Block*.

12 Select *OK*.

Position these at the end of the URL filter list so that any exemptions or blocks before that are still effective.

Both of these filter entries are required. If you only enter the second one, the FortiGate unit will also catch a URL lookup as they both behave in a similar fashion after the URL is resolved to an IP. The first entry is needed to break out of the URL filter and allow the web site before it does the second check if they entered text.

ICMP packet processing

ICMP messages are used to relay feedback to the traffic source that the destination IP is not reachable. ICMP message types are:

- ICMP_ECHO
- ICMP_TIMESTAMP
- ICMP_INFO_REQUEST
- ICMP_ADDRESS

For ICMP error messages, only those reporting an error for an existing session can pass through the firewall. The security policy will allow traffic to be routed, forwarded or denied. If allowed, the ICMP packets will start a new session. Only ICMP error messages of a corresponding security policy is available will be sent back to the source. Otherwise, the packet is dropped. That is, only ICMP packets for a corresponding security policy can traverse the FortiGate unit.

Common error messages include:

- destination unreachable messages
- time exceeded messages
- redirect messages

For example, a security policy that allows TFTP traffic through the FortiGate unit. User1 (192.168.21.12) attempts to connect to the TFTP server (10.11.100.1), however, the UDP port 69 has not been opened on the server. The corresponding sniffer trace occurs:

```
diagnose sniffer packet any "host 10.11.100.1 or icmp 4"
3.677808 internal in 192.168.21.12.1262 -> 10.11.100.1.69: udp 20
3.677960 wan1 out 192.168.21.12.1262 -> 10.11.100.1.69: udp 20
3.678465 wan1 in 10.11.100.1.132 -> 192.168.21.12: icmp:
10.11.100.1 udp port 69 unreachable
3.678519 internal out 10.11.100.1 -> 192.168.21.12: icmp:
192.168.182.132 udp port 69 unreachable
```

Adding NAT security policies in Transparent mode

Similar to operating in NAT mode, when operating a FortiGate unit in Transparent mode you can add security policies and:

- Enable NAT to translate the source addresses of packets as they pass through the FortiGate unit.
- Add virtual IPs to translate destination addresses of packets as they pass through the FortiGate unit.

- Add IP pools as required for source address translation

For NAT firewall policies to work in NAT mode you must have two interfaces on two different networks with two different subnet addresses. Then you can create firewall policies to translate source or destination addresses for packets as they are relayed by the FortiGate unit from one interface to the other.

A FortiGate unit operating in Transparent mode normally has only one IP address, the management IP. To support NAT in Transparent mode, you can add a second management IP. These two management IPs must be on different subnets. When you add two management IP addresses, all FortiGate unit network interfaces will respond to connections to both of these IP addresses.

In the example shown in [Figure 11](#), all of the PCs on the internal network (subnet address 192.168.1.0/24) are configured with 192.168.1.99 as their default route. One of the management IPs of the FortiGate unit is set to 192.168.1.99. This configuration results in a typical NAT mode firewall. When a PC on the internal network attempts to connect to the Internet, the PC's default route sends packets destined for the Internet to the FortiGate unit internal interface. Similarly on the DMZ network (subnet address 10.1.1.0/24) all of the PCs have a default route of 10.1.1.99.

This example describes adding an internal to WAN1 security policy to relay these packets from the internal interface out the WAN1 interface to the Internet. Because the WAN1 interface does not have an IP address of its own, you must add an IP pool to the WAN1 interface that translates the source addresses of the outgoing packets to an IP address on the network connected to the wan1 interface.


The example describes adding an IP pool with a single IP address of 10.1.1.201. So all packets sent by a PC on the internal network that are accepted by the Internal to WAN1 policy leave the WAN1 interface with their source address translated to 10.1.1.201. These packets can now travel across the Internet to their destination. Reply packets return to the WAN1 interface because they have a destination address of 10.1.1.201. The Internal to WAN1 NAT policy translates the destination address of these return packets to the IP address of the originating PC and sends them out the internal interface to the originating PC.

Use the following steps to configure NAT in Transparent mode

- Add two management IPs
- Add an IP pool to the WAN1 interface
- Add an Internal to WAN1 security policy



You can add the security policy from the web-based manager and then use the CLI to enable NAT and add the IP pool.

Figure 11: Example NAT in Transparent mode configuration**To add a source address translation NAT policy in Transparent mode**

- 1 Enter the following command to add two management IPs.

The second management IP is the default gateway for the internal network.

```
config system settings
  set manageip 10.1.1.99/24 192.168.1.99/24
end
```

- 2 Enter the following command to add an IP pool to the WAN1 interface:

```
config firewall ippool
  edit nat-out
    set interface "wan1"
    set startip 10.1.1.201
    set endip 10.1.1.201
  end
```


- 3 Enter the following command to add an Internal to WAN1 security policy with NAT enabled that also includes an IP pool:

```
config security policy
  edit 1
    set srcintf "internal"
    set dstintf "wan1"
    set scraddr "all"
    set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ANY"
    set nat enable
    set ippool enable
    set poolname nat-out
  end
```


Adding a static NAT virtual IP for a single IP address and port

In this example, the wan1 interface of the FortiGate unit is connected to the Internet and the Internal interface is connected to the DMZ network. The IP address 192.168.37.4 on port 80 on the Internet is mapped to 10.10.10.42 on port 8000 on the private network. Attempts to communicate with 192.168.37.4 from the Internet are translated and sent to 10.10.10.42 by the FortiGate unit. The computers on the Internet are unaware of this translation and see a single computer at 192.168.37.4 rather than a FortiGate unit with a private network behind it.

Figure 12: Static NAT virtual IP for a single IP address example



To add a static NAT virtual IP for a single IP address and port - web-based manager

- 1 Go to *Firewall Objects > Virtual IP > Virtual IP*.
- 2 Select *Create New*.
- 3 Complete the following and select *OK*.

Name	static_NAT
External Interface	wan1
Type	Static NAT
External IP Address/Range	192.168.37.4.
Mapped IP Address/Range	10.10.10.42
Port Forwarding	Selected
Protocol	TCP
External Service Port	80
Map to Port	8000

To add a static NAT virtual IP for a single IP address and port - CLI

```
config firewall vip
  edit static_NAT
    set extintf wan1
    set type static-nat
    set extip 192.168.37.4
    set mappedip 10.10.10.42
    set portforward enable
    set extport 80
    set mappedport 8000
  end
```

Add a external to dmz1 security policy that uses the virtual IP so that when users on the Internet attempt to connect to the web server IP address packets pass through the FortiGate unit from the external interface to the dmz1 interface. The virtual IP translates the destination address of these packets from the external IP to the DMZ network IP address of the web server.

To add a static NAT virtual IP for a single IP address to a security policy - web-based manager

- 1 Go to *Policy > Policy > Policy* and select *Create New*.
- 2 Complete the following:

Source Interface/Zone	wan1
Source Address	All
Destination Interface/Zone	Internal
Destination Address	static_nat
Schedule	always
Service	HTTP
Action	ACCEPT

- 3 Select *NAT*.
- 4 Select *OK*.

To add a static NAT virtual IP for a single IP address to a security policy - CLI

```
config security policy
  edit 1
    set srcintf wan1
    set dstintf internal
    set srcaddr all
    set dstaddr static_nat
    set action accept
    set schedule always
    set service ANY
    set nat enable
  end
```


Double NAT: combining IP pool with virtual IP

In this example, a combination of virtual IPs, IP pools and security policies will allow the local users to access the servers on the DMZ. The example uses a fixed port and IP pool to allow more than one user connection while using virtual IP to translate the destination port from 8080 to 80. The security policy uses both the IP pool and the virtual IP for double IP and/or port translation.

For this example:

- Users in the 10.1.1.0/24 subnet use port 8080 to access server 172.20.120.1.
- The server's listening port is 80.
- Fixed ports must be used.

Figure 13: Double NAT



To create an IP pool - web-based manager

- 1 Go to *Firewall Objects > Virtual IP > IP Pool*.
- 2 Select *Create New*.
- 3 Enter the *Name* pool-1.
- 4 Enter the *IP Range/Subnet* 10.1.3.1-10.1.3.254.
- 5 Select *OK*.

To create an IP pool - CLI

```
config firewall ippool
  edit pool-1
    set startip 10.1.3.1
    set endip 10.1.3.254
  end
```

Next, create the virtual IP with port translation to translate the user internal IP used by the network users to the DMZ port and IP address of the server.

To create a Virtual IP with port translation - web-based manager

- 1 Go to *Firewall Objects > Virtual IP > Virtual IP*.
- 2 Select *Create New*.

3 Enter the following information and select **OK**.

Name	server-1
External Interface	Internal
Type	Static NAT
External IP Address/Range	172.20.120.1
	Note: This address is the same as the server address.
Mapped IP Address/Range	172.20.120.1
Port Forwarding	Enable
Protocol	TCP
External Service Port	8080
Map to Port	80

To create a Virtual IP with port translation - CLI

```
config firewall vip
  edit server-1
    set extintf internal
    set type static-nat
    set extip 172.20.120.1
    set mappedip 172.20.120.1
    set portforward enable
    set extport 80
    set mappedport 8080
  end
```

Add an internal to DMZ security policy that uses the virtual IP to translate the destination port number and the IP pool to translate the source addresses.

To create the security policy - web-based manager

- 1 Go to *Policy > Policy > Policy* and select *Create New*.
- 2 Complete the following and select **OK**:

Source Interface/Zone	internal
Source Address	all
Destination Interface/Zone	dmz
Destination Address	server-1
Schedule	always
Service	HTTP
Action	ACCEPT
NAT	Select
Dynamic IP Pool	Select, and select the <i>pool-1</i> IP pool.

To create the security policy - CLI

```
config security policy
```


```
edit 1
  set srcintf internal
  set dstintf dmz1
  set srcaddr all
  set dstaddr server-1
  set action accept
  set schedule always
  set service HTTP
  set nat enable
  set ippool enable
  set poolname pool-1
end
```

Using VIP range for Source NAT (SNAT) and static 1-to-1 mapping

VIP addresses are typically used to map external (public) to internal (private) IP addresses for Destination NAT (DNAT).

This example shows how to use VIP ranges to perform source NAT (SNAT) with a static 1-to-1 mapping from internal to external IP addresses. This is similar to using an IP pool with the advantage of having predictable and static 1-to-1 address mapping.

Figure 14: Network diagram



This example will associate each internal IP address to one external IP address for the Source NAT (SNAT) translation.

Using the diagram above, the translations will look like the following:

Traffic from Source IP Translated to Source IP (SNAT)

10.10.10.42	192.168.37.4
10.10.10.43	192.168.37.5

```
...
10.10.10.46      192.168.37.8
```

First, configure the virtual IP.

To configure the virtual IP - web-based manager

- 1 Go to *Firewall Objects > Virtual IP > Virtual IP* and select *Create New*.
- 2 Enter the *Name* of `Static_NAT_1to1`.
- 3 Select the *External Interface* of *port 1* from the drop-down list.
- 4 Enter the *External IP Address* of `192.168.37.4`.
- 5 Enter the *Mapped IP Address* range of `10.10.10.42 to 10.10.10.46`.
- 6 Select OK.

To configure the virtual IP - CLI

```
config firewall vip
  edit "Static_NAT_1to1"
    set extip 192.168.37.4
    set extintf "port1"
    set mappedip 10.10.10.42-10.10.10.46
  next
end
```

Next, configure the firewall policies. Even if no connection needs to be initiated from external to internal, a second security policy number is required to activate the VIP range. Otherwise the IP address of the physical interface is used for NAT. In this example it is set as a "DENY" security policy for security purpose.

To configure the firewall policies - web-based manager

- 1 Go to *Policy > Policy > Policy* and select *Create New*.
- 2 Complete the following and select OK:

Source Interface/Zone	port2
Source Address	all
Destination Interface/Zone	port1
Destination Address	all
Schedule	always
Service	ANY
Action	ACCEPT
NAT	Select

- 3 Complete the following and select OK:

Source Interface/Zone	port 1
Source Address	all
Destination Interface/Zone	port 2
Destination Address	Static_NAT_1to1

Schedule	always
Service	ALL
Action	deny
Comments	Used to activate static Source NAT 1-to-1

To configure the firewall policies - CLI

```

config firewall policy
  edit 1
    set srcintf port2
    set dstintf port1
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
    set nat enable
  next
  edit 2
    set srcintf port1
    set dstintf port2
    set srcaddr all
    set dstaddr Static_NAT_1to1
    set schedule always
    set service ANY
    set action deny
    set comments (Used to activate static Source NAT 1-to-1)
  next
end
end

```

Traffic shaping and per-IP traffic shaping

Traffic shaping helps to optimize traffic flow through the FortiGate unit, and per-IP traffic shaping does much the same, however, it applies traffic shaping per IP address instead of per policy or per shaper. Traffic shaping, when included in a security policy, controls the bandwidth available to the policy, and sets the priority of the traffic processed by the policy. Traffic shaping makes it possible to control which policies have the highest priority when large amounts of data are moving through the FortiGate unit. For example, the policy for the corporate web server might be given higher priority than the policies for most employee's computers. An employee who needs extra high speed Internet access could have a special outgoing policy set up with higher bandwidth.

Traffic shaping is available for security policies whose Action is ACCEPT, IPSEC, or SSL VPN. It is also available for all supported services, including H.323, TCP, UDP, ICMP, and ESP.

Traffic shaping is used to improve the quality of bandwidth-intensive and sensitive traffic; it also cannot increase the total amount of bandwidth available. The bandwidth available for traffic set in a traffic shaper is used to control data sessions for traffic in both directions.

For more information about traffic shaping, see the Traffic Shaping chapter in the [FortiOS Handbook](#).

Endpoint Security

Endpoint security enforces the use of the FortiClient End Point Security (FortiClient and FortiClient Lite) application on your network. It can also allow or deny endpoints access to the network based on the application installed on them.

By applying endpoint security to a security policy, you can enforce this type of security on your network. FortiClient enforcement can check that the endpoint is running the most recent version of the FortiClient application, that the antivirus signatures are up-to-date, and that the firewall is enabled. An endpoint is usually often a single PC with a single IP address being used to access network services through a FortiGate unit.

With endpoint security enabled on a policy, traffic that attempts to pass through, the FortiGate unit runs compliance checks on the originating host on the source interface. Non-compliant endpoints are blocked. If someone is browsing the web, the endpoints are redirected to a web portal which explains the non-compliance and provides a link to download the FortiClient application installer. The web portal is already installed on the FortiGate unit, as a replacement message, which you can modify if required.

Endpoint Security requires that all hosts using the security policy have the FortiClient Endpoint Security agent installed. Currently, FortiClient Endpoint Security is available for Microsoft Windows 2000 and later only.

For more information about endpoint security, see the UTM chapter in the [FortiOS Handbook](#).

Logging traffic

When you enable logging on a security policy, the FortiGate unit records the scanning process activity that occurs, as well as whether the FortiGate unit allowed or denied the traffic according to the rules stated in the security policy. This information can provide insight into whether a security policy is working properly, as well as if there needs to be any modifications to the security policy, such as adding traffic shaping for better traffic performance.

Traffic is logged in the traffic log file and provides detailed information that you may not think you need, but do. For example, the traffic log can have information about an application used (web: HTTP.Image), and whether or not the packet was SNAT or DNAT translated. The following is an example of a traffic log message.

```
2011-04-13 05:23:47 log_id=4 type=traffic subtype=other pri=notice
vd=root status="start" src="10.41.101.20" srcname="10.41.101.20"
src_port=58115 dst="172.20.120.100" dstname="172.20.120.100"
dst_country="N/A" dst_port=137 tran_ip="N/A" tran_port=0
tran_sip="10.31.101.41" tran_sport=58115 service="137/udp"
proto=17 app_type="N/A" duration=0 rule=1 policyid=1 sent=0 rcvd=0
shaper_drop_sent=0 shaper_drop_rcvd=0 perip_drop=0
src_int="internal" dst_int="wan1" SN=97404 app="N/A" app_cat="N/A"
carrier_ep="N/A"
```

If you want to know more about logging, see the Logging and Reporting chapter in the [FortiOS Handbook](#). If you want to know more about traffic log messages, see the [FortiGate Log Message Reference](#).

Quality of Service (QoS)

The Quality of Service (QoS) feature is an advanced firewall component that applies bandwidth limits and prioritization to traffic. QoS is the capability of the network to adjust some quality aspects for selected flows within your overall network traffic, and may include such techniques as priority-based queuing and traffic policing.

QoS can be implemented for services that include H.323, TCP, UDP, ICMP, and ESP. QoS uses the following techniques:

Traffic policing	Drops packets that do not conform to bandwidth limitations
Traffic shaping	This helps to ensure that the traffic may consume bandwidth at least at the guaranteed rate by assigning a greater priority queue if the guarantee is not being met. Traffic shaping also ensures that the traffic cannot consume bandwidth greater than the maximum at any given instant in time. Flows that are greater than the maximum rate are subject to traffic policing.
Queuing	This transmits packets in order of their assigned priority queue for that physical interface. All traffic in a higher priority traffic queue must be completely transmitted before traffic in lower priority queues will be transmitted.

QoS can be helpful for organizations that are trying to manage their voice and streaming multi-media traffic, which can rapidly consume bandwidth. Both voice and streaming multi-media are sensitive to latency.

For additional information about QoS, see the Traffic Shaping chapter in the [FortiOS Handbook](#).

Identity-based security policies

Identity-based security policies, also known as authentication policies, match traffic that requires a supported authentication protocol to trigger the firewall authentication challenge and successfully authenticate network users. Network users authentication can occur using HTTP, HTTPS, FTP, and Telnet protocols as well as through automatic login using NTLM and FSSO, to bypass user intervention.

Identity-based security policies are usually configured for IPsec or SSL VPN traffic since this type of traffic usually requires authentication from network users.

When configuring identity-based policies, you can use schedules to limit network users authentication sessions. For example, example.com has a schedule policy to use P2P applications between noon and 1:00 pm, and a user authentication timeout of 30 minutes. When a user logs in at 12:15 pm, their authentication time logs them off at 12:45 (30 minutes later). You can configure this type of authentication by using the `schedule-timeout` field in the `config firewall policy` command in the CLI.

Identity-based policy positioning

With identity-based security policies, positioning is extremely important. For a typical security policy, the FortiGate unit matches the source, destination and service of the policy. If matched, it acts on that policy. If not, the FortiGate unit moves to the next policy.

With identity-based policies, once the FortiGate unit matches the source and destination addresses, it processes the identity sub-rules for the user groups and services. That is, it acts on the authentication and completes the remainder of that policy and goes no further in the policy list.

The way identity based policies work is that once src/dest are matched, it will process the identity based sub-rules (for lack of a better term) around the user groups and services. It will never process the rest of your rulebase. For this reason, unique security policies should be placed **before** an identity-based policy.

For example, consider the following policies:

Seq. No.	Source	Destination	Schedule	Service	Action	Status	Authentication
1	all	all	always	DNS	ACCEPT	<input checked="" type="checkbox"/>	
2	all	all	always	HTTP HTTPS	ACCEPT	<input checked="" type="checkbox"/>	FSAE_Guest_Users
3	all	all	always	ANY	DENY	Implicit	

DNS traffic goes through successfully as does any HTTP traffic after being authenticated. However, if there was FTP traffic, it would not get through. As the FortiGate unit processes FTP traffic, it skips rule one since it's matching the source, destination and service. When it moves to rule two it matches the source and destination, it determines there is a match and, sees there are also processes the group/service rules, which requires authentication and acts on those rules. Once satisfied, the FortiGate unit will never go to rule three.

In this situation, where you would want FTP traffic to traverse the FortiGate unit, create a security policy specific to the services you require and place it above the authentication policy.

Identity-based sub-policies

When adding authentication to a security policy, you can add multiple authentication rules, or sub-policies. Within these policies you can include additional UTM profiles, traffic shaping and so on, to take affect on the selected services.

Figure 15: Authentication sub-policies

☒ Enable Identity Based Policy

Add

Rule ID	User Group	Service	Schedule	UTM	Traffic Shaping	Logging	
1	FSAE_Guest_Users	HTTP,HTTPS	always	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	
2	FSAE_Guest_Users	FTP	always	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	

☒ Firewall
 ☒ Directory Service(FSAE)
 ☐ NTLM Authentication

These sub-policies work on the same principle as normal security policies, that is, top down until the criteria has been met. As such, if there is no matching policy within the list, the packet can still be dropped even after authentication is successful.



Index

A

- accept, 46
- adding configuring defining
 - deny security policy, 52
 - how firewall components create a FortiGate firewall, 12
 - how packets flow, 14
 - how to apply VLANs and zones to security policies, 23
 - how to arrange security policies, 49
 - how to create a basic security policy for Internet access, 55
- interfaces and zones, 23
- ipv4 tunneling, 65
- ipv6, dual stack routing, 64
- remotely connecting to an IPv6 over the Internet, 65
- adding, configuring defining
 - central NAT table, 94
 - tags, 83
- address
 - CIDR format, 24
 - FDQN, 31
 - geography-based, 28
 - groups, 32
 - IP pool, 26
 - IP range, 25
 - IPv6, 63
 - matching, IP pool, 27
- addresses
 - ipv6, 63
- addresses, firewall, 24
- AFS3, advanced file security encrypted file
 - AFS3, 34
- AH, predefined service, 34
- ANY
 - service, 34
- AOL
 - service, 34
- arranging security policies, 49

B

- BGP
 - service, 34
- BGP, IPv6, 79
- blocking
 - http access by ip, 101
 - port 25, 98

C

- central NAT, 93
- central NAT table
 - configuring, 94
- Classless Inter-Domain Routing (CIDR), 66
- column settings, security policies, 46

- comments
 - firewall policy, 83
- custom services, 33
- CVSPSERVER, concurrent versions system proxy server, 34, 35

D

- DCE-RPC
 - firewall service, 34, 35
- default
 - password, 9
- deny, 46
- deny policy, 52
- destination
 - firewall policy, 80
- details, security policies, 46
- DHCP (Dynamic Host Configuration Protocol)
 - service, 34, 35
- DHCP6
 - service, 34, 35
- DNS
 - service, 34, 35
 - TTL, 31
- double NAT example, 107
- dual stack routing, ipv6, 64
- dynamic IP pool, 81

E

- ESP
 - service, 34, 35

F

- FDQN, 31
- FINGER
 - service, 34, 35
- firewall
 - applying VLANs and zones to security policies, 23
 - central NAT table, 94
 - dynamic IP pool, 81
 - how firewall components create a FortiGate firewall, 12
 - interfaces and zones, 23
 - ipv6, 63
 - option, any, 79
 - predefined services, 33
 - source interface, 79
 - what is it, 11
 - wildcard addresses, 29
- firewall address, 24
- firewall addresses
 - ipv6, 63
 - wildcard address, 29

- firewall policies
 - adding NAT policies to transparent mode, 102
- firewall policy
 - comments, 83
 - destination, 80
 - log traffic, 80
 - schedule, 80
 - service, 80
 - traffic shaping, 82

firewall service

- AFS3, 34
- AH, 34
- ANY, 34
- AOL, 34
- BGP, 34
- CVSPSERVER, 34, 35
- DCE-RPC, 34, 35
- DHCP, 34, 35
- DHCP6, 34, 35
- DNS, 34, 35
- ESP, 34, 35
- FINGER, 34, 35
- FTP, 34, 35
- FTP_GET, 34, 35
- H323, 36
- HTTP, 36
- HTTPS, 36
- ICMP_ANY, 36
- IKE, 36
- IMAP, 36
- INFO_ADDRESS, 36
- INFO_REQUEST, 36
- Internet-Locator-Service, 36
- IRC, 36
- L2TP, 36
- LDAP, 36
- MGCP, 36
- MS-SQL, 37
- MYSQL, 37
- NetMeeting, 37
- NFS, 37
- NNTP, 37
- NTP, 37
- ONC-RPC, 37
- OSPF, 37
- PC-Anywhere, 37
- PING, 37
- PING6, 37
- POP3, 37
- PPTP, 38
- QUAKE, 38
- RAUDIO, 38
- REXEC, 38
- RIP, 38
- RLOGIN, 38
- RSH, 38
- RTSP, 38
- SAMBA, 38
- SCCP, 38
- SIP, 39
- SIP-MSNmessenger, 39
- SMTP, 39
- SNMP, 39
- SOCKS, 39
- SQUID, 39
- SSH, 39
- SYSLOG, 39
- TALK, 39
- TCP, 39
- TELNET, 39
- TFTP, 40
- TIMESTAMP, 40
- UDP, 40

- UUCP, 40

- VDOLIVE, 40

- viewing predefined list, 33

- VNC, 40

- WAIS, 40

- WINFRAME, 40

- WINS, 40

- X-WINDOWS, 40

- fixed ports, IP pools, 27

- FortiGate firewall

- creating, 12

- FortiGuard

- Antispam, 9

- Antivirus, 9

- FortiOS

- ipv6, 64

- FTP

- service, 34, 35

- FTP_GET

- service, 34, 35

G

- geography-based addressing, 28

- groups, addressing, 32

H

- H323

- service, 36

- how to allow DNS queries to only one DNS server, 52

- how to apply VLANs and zones to security policies, 23

- how to arrange policies, 49

- how to create basic security policy for Internet access, 55

- how to test basic security, 55

- how to use match-vip, 33

- how to use UTM profiles to monitor and protect your network, 42

- HTTP

- service, 36

- http

- blocking, 101

- HTTPS

- service, 36

I

- ICMP processing, 102

- ICMP_ANY

- service, 36

- identity-based policy, 50

- position, 113

- IEEE 1394 (FireWire), 67

- IKE

- service, 36

- IMAP

- service, 36

- INFO_ADDRESS

- service, 36

- INFO_REQUEST

- service, 36

- interfaces, 23
 - ANY, ANY interface option, 47
 - source, firewall policy, 79
- Internet-Locator-Service
 - service, 36
- IP pool, 26
 - address matching, 27
 - policies and fixed ports, 27
- IP range, 25
- IPsec, 46
- ipv4 tunneling configuration, ipv6, 65
- IPv6, 63
 - dual stack, 73
 - dynamic routing, 79
 - interfaces, 77
 - IPsec certificate configuration, 85
 - IPSec configuration, 85
 - IPSec phase 1 configuration, 85
 - IPsec phase 2 configuration, 86
 - IPsec routing configuration, 86
 - IPsec security policy configuration, 86
 - Neighbor Discovery (ND), 72
 - security policies, 79
 - static routing, 78
 - troubleshooting, 86
 - tunnel provider example, 74
 - tunneling, 73
- ipv6, 63
 - dual stack routing configuration, 64
 - ipv4 tunneling configuration, 65
 - remotely connecting over the Internet, 65
- ipv6 in FortiOS, 64
- IRC
 - service, 36

L

- L2TP
 - service, 36
- LDAP
 - service, 36
- life of a packet, 13
- local-in policy, 54
- log
 - traffic, firewall policy, 80

M

- M3UA, 96
- match-vip, 32
 - how to, 33
- Maximum Transmission Unit (MTU), 66
- Message Transfer Part 3, 96
- MGCP
 - service, 36
- mode
 - operation, 9
- MS-SQL
 - service, 37
- MTP3 User Adaptation Layer, 96
- MYSQL
 - service, 37

N

- NAT, 93
- Neighbor
 - Advertisement, 73
 - Solicitation, 73
- netmask
 - wildcard firewall addresses, 29
- NetMeeting
 - service, 37
- NFS
 - service, 37
- NNTP
 - service, 37
- NTP
 - service, 37

O

- ONC-RPC
 - service, 37
- operation mode, 9
- OSPF
 - IPv6, 79
 - service, 37

P

- packet
 - ICMP, 102
 - life of, 13
- packet flow, 14
- packets
 - flow, 14
- password
 - administrator, 9
- PC-Anywhere
 - service, 37
- PING
 - service, 37
- PING6
 - firewall service, 37
- policies, 46
 - column settings, 46
 - expiry, 41
 - ICMP packets, 102
 - identity-based, 50
 - NAT to transparent mode, 102
 - order, 47
 - timeout, 41
 - viewing, 47
- policy
 - comments, 83
 - local-in, 54
 - log traffic, 80
 - schedule, 80
 - service, 80
 - traffic shaping, 82
- policy 0, 53
- POP3
 - service, 37
- port
 - blocking port 25, 98

- ports
 - services, 33
- position
 - identity-based policy, 113
- PPTP
 - service, 38
- predefined services, 33
- protocol
 - service, 34
- PSTN, 95
- Public Switched Telephone Network
 - See PSTN, 95

Q

- QUAKE
 - service, 38

R

- RAUDIO
 - service, 38
- Redirect message, 73
- remotely connecting to IPv6 over the Internet, 65
- REXEC
 - firewall service, 38
- RFC
 - 2071, 66
 - 2080, 79
 - 2185, 73
 - 2545, 79
 - 2740, 79
 - 2858, 79
 - 2893, 73
- RIP
 - service, 38
- RIP, IPv6, 79
- RLOGIN
 - service, 38

- Router Solicitation message, 73
- RSH
 - firewall service, 38
- RTSP
 - firewall service, 38

S

- SAMBA
 - service, 38
- SCCP
 - firewall service, 38
- schedule
 - firewall policy, 80
 - timeout, 41
- schedules
 - expiry, 41
 - group, 41
 - one time, 41
 - recurring, 41
- schedule-timeout command, 41
- security policies, 53
 - accept, 46
 - column settings, 46
 - deny, 46
 - deny policy, 52
 - how to apply VLANs and zones, 23
 - how to arrange, 49
 - ICMP packets, 102
 - identity-based, 50
 - IPsec, 46
 - policy order, 47
 - ssl-vpn policies, 46
 - viewing, 47
- security policy
 - how to allow Internet access, 55
 - local-in, 54
 - verifying traffic is hitting a policy, 56

- service
 - AH, 34
 - ANY, 34
 - AOL, 34
 - BGP, 34
 - CVSPSERVER, 34, 35
 - DCE-RPC, 34, 35
 - DHCP, 34, 35
 - DHCP6, 34, 35
 - DNS, 34, 35
 - ESP, 34, 35
 - FINGER, 34, 35
 - firewall policy, 80
 - FTP, 34, 35
 - FTP_GET, 34, 35
 - H323, 36
 - HTTPS, 36
 - ICMP_ANY, 36
 - IKE, 36
 - IMAP, 36
 - INFO_ADDRESS, 36
 - INFO_REQUEST, 36
 - Internet-Locator-Service, 36
 - IRC, 36
 - L2TP, 36
 - LDAP, 36
 - MGCP, 36
 - MS-SQL, 37
 - MYSQL, 37
 - NetMeeting, 37
 - NFS, 37
 - NNTP, 37
 - NTP, 37
 - ONC-RPC, 37
 - OSPF, 37
 - PC-Anywhere, 37
 - PING, 37
 - PING6, 37
 - POP3, 37
 - PPTP, 38
 - predefined, 33
 - QUAKE, 38
 - RAUDIO, 38
 - REXEC, 38
 - RIP, 38
 - RLOGIN, 38
 - RSH, 38
 - RTSP, 38
 - SAMBA, 38
 - SCCP, 38
 - service name, 34
 - SIP, 39
 - SIP-MSNmessenger, 39
 - SMTP, 39
 - SNMP, 39
 - SOCKS, 39
 - SQUID, 39
 - SSH, 39
 - SYSLOG, 39
 - TALK, 39
 - TCP, 39
 - TELNET, 39
 - TFTP, 40
 - TIMESTAMP, 40
 - UDP, 40
 - UUCP, 40
 - VDOLIVE, 40
 - VNC, 40
 - WAIS, 40
 - WINFRAME, 40
 - WINS, 40
 - X-WINDOWS, 40
 - services, 33
 - custom, 33
 - list, 33
 - Simple Internet Transition (SIT), 84
 - SIP
 - service, 39
 - SIP-MSNmessenger
 - service, 39
 - SMTP
 - service, 39
 - smtp traffic, 98
 - SNMP
 - service, 39
 - SOCKS
 - service, 39
 - SQUID
 - service, 39
 - SS7, 96
 - SSH
 - service, 39
 - SSL
 - service definition, 36, 37
 - ssl-vpn, 46
 - SYSLOG
 - service, 39
- ## T
- tags
 - adding tags, 83
 - applying tags, 83
 - TALK
 - service, 39
 - TCP
 - service, 39
 - TELNET
 - service, 39
 - testing a basic security policy, 55
 - TFTP
 - service, 40
 - TIMESTAMP
 - service, 40
 - traffic shaping
 - firewall policy, 82
 - transparent mode
 - adding NAT policies, 102
 - tunnel provider, IPv6, 74
 - tunneling, IPv6, 73
- ## U
- UDP service, 40
 - understanding firewall addresses, 24
 - using UTM profiles to monitor and protect your network, 42

UTM
 profiles, 42
UUCP
 service, 40

V

VDOLIVE
 service, 40
verifying traffic is hitting a policy, 56
viewing
 firewall predefined service list, 33
viewing security policies, 47
vip, 32
vip, grouping, 32
vip, match-vip, 32
virtual ip addresses, 32
VNC
 service, 40

W

WAIS
 service, 40
wildcard
 firewall addresses, 29
wildcard addresses, 29
WINFRAME
 service, 40
WINS
 service, 40

X

X-WINDOWS
 service, 40

Z

zones, 23

